University of Groningen

Stochastic collusion and the power law of learning: a general reinforcement learning model of cooperation

Flache, A.

Published in:
Journal of Conflict Resolution

DOI:
[10.1177/002200202236167](https://doi.org/10.1177/002200202236167)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2002

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Flache, A. (2002). Stochastic collusion and the power law of learning: a general reinforcement learning model of cooperation. *Journal of Conflict Resolution*, 46(5), 629-653.
<https://doi.org/10.1177/002200202236167>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Stochastic Collusion and the Power Law of Learning

A GENERAL REINFORCEMENT LEARNING MODEL OF COOPERATION

ANDREAS FLACHE

*Department of Sociology
University of Groningen, the Netherlands*

MICHAEL W. MACY

*Department of Sociology
Cornell University*

Concerns about models of cultural adaptation as analogs of genetic selection have led cognitive game theorists to explore learning-theoretic specifications. Two prominent examples, the Bush-Mosteller stochastic learning model and the Roth-Erev payoff-matching model, are aligned and integrated as special cases of a general reinforcement learning model. Both models predict *stochastic collusion* as a backward-looking solution to the problem of cooperation in social dilemmas based on a random walk into a self-reinforcing cooperative equilibrium. The integration uncovers hidden assumptions that constrain the generality of the theoretical derivations. Specifically, Roth and Erev assume a “power law of learning”—the curious but plausible tendency for learning to diminish with success and intensify with failure. Computer simulation is used to explore the effects on stochastic collusion in three social dilemma games. The analysis shows how the integration of alternative models can uncover underlying principles and lead to a more general theory.

FROM EVOLUTIONARY TO COGNITIVE GAME THEORY

The escalation of conflict between neighbors, friends, business associates, ethnic groups, and states strikes observers as tragically misguided and irrational, yet these same observers may find themselves trapped in similar dynamics without intending to be. Why do some relationships proceed cooperatively whereas others descend into mutually detrimental discord, rivalry, dishonesty, or injury?

AUTHORS' NOTE: This research was made possible by a fellowship of the Royal Netherlands Academy of Arts and Sciences to the first author and the U.S. National Science Foundation (SES-0079381) to the second author. We also thank Yoshimichi Sato, Douglas Heckathorn, and other members of the Janus Workshop on Backward-Looking Rationality at Cornell University for their helpful suggestions. Please direct correspondence to the first author. The computer routines used in this article are available at <http://www.yale.edu/unsy/jcr/jcrdata.htm>.

JOURNAL OF CONFLICT RESOLUTION, Vol. 46 No. 5, October 2002 629-653
DOI: 10.1177/0022002202236167
© 2002 Sage Publications

Game theory has formalized the problem of cooperation at the most elementary level as a mixed-motive, two-person game with two choices: cooperate or defect. These choices intersect at four possible outcomes, abbreviated as *CC*, *CD*, *DD*, and *DC*. Each outcome has an associated payoff: *R* (reward), *S* (sucker), *P* (punishment), and *T* (temptation). Using these payoffs, we define a two-person social dilemma¹ as any ordering of these payoffs such that mutual cooperation is Pareto optimal yet may be undermined by the temptation to cheat (if $T > R$) or the fear of being cheated (if $P > S$) or both. In the game of Stag Hunt, the problem is fear but not greed ($R > T > P > S$); and in the game of Chicken, the problem is greed but not fear ($T > R > S > P$). The problem is most challenging when both fear and greed are present, that is, when $T > R$ and $P > S$. Given the assumption that $R > P$, there is only one way this can happen—if $T > R > P > S$, the celebrated game of Prisoner's Dilemma (PD).

The Nash equilibrium²—the main solution concept in analytical game theory—predicts mutual defection in PD, unilateral defection in Chicken, and either mutual cooperation or mutual defection in Stag Hunt. However, Nash cannot make precise predictions about the selection of supergame equilibria, that is, about the outcome of ongoing, mixed-motive games. Nor can it tell us much about the dynamics by which a population of players can move from one equilibrium to another.

These limitations, along with concerns about the cognitive demands of forward-looking rationality (Dawes and Thaler 1988; Weibull 1998; Fudenberg and Levine 1998), have led game theorists to explore backward-looking alternatives based on evolution and learning. Evolutionary models test the ability of conditionally cooperative strategies to survive and reproduce in competition with predators (Maynard-Smith 1982). Biological models have also been extended to military and economic games in which losers are physically eliminated or swallowed up and to cultural games in which winners are more likely to be imitated (Axelrod 1984; Weibull 1998). However, despite rapid growth in the science of mimetics (Dawkins 1989; Dennett 1995; Blackmore 1999), skeptics charge that genetic learning may be a misleading template for models of adaptation at the cognitive level (Aunger 2001).

The need for a cognitive alternative to evolutionary game theory is reflected in a growing number of formal learning-theoretic models of cooperative behavior (Macy 1991; Roth and Erev 1995; Fudenberg and Levine 1998; Peyton-Young 1998; Cohen, Riolo, and Axelrod 2001). In learning, positive outcomes increase the probability that the associated behavior will be repeated, whereas negative outcomes reduce it. The process closely parallels evolutionary selection, in which positive outcomes increase a strategy's chances for survival and reproduction whereas negative outcomes reduce it (Rapoport 1994; Axelrod and Cohen 2000). However, this isomorphism need not imply that adaptive actors will learn the strategies favored by evolutionary selection

1. Dawes (1980) defined a social dilemma as an *n*-person, mixed-motive game in which a Pareto-deficient outcome is the *only* individually rational course of action, as in Prisoner's Dilemma. Raub (1988; see also Liebrand 1983) proposed a less restrictive criterion—that there is at least one Pareto-deficient individually rational outcome of the game, as in Stag Hunt and Chicken as well as in Prisoner's Dilemma. Unlike Dawes's definition, Raub's criterion does not exclude the possibility that rational behavior may also result in Pareto-efficient outcomes. Our usage follows Raub.

2. In a Nash equilibrium, no one has an incentive to unilaterally change strategies, given the expected utility of the alternatives.

Figure 1: Schematic Model of Reinforcement Learning

NOTE: C = cooperate; D = defect; R = reward; S = sucker; P = punishment; T = temptation.

pressures (Fudenberg and Levine 1998; Schlag and Pollock 1999). In evolution, strategies compete *between* the individuals that carry them, not *within* them. That is, evolutionary models explore changes in the global frequency distribution of strategies across a population, whereas learning operates on the local probability distribution of strategies within the repertoire of each individual member. Put differently, learning models provide a microfoundation for extragenetic evolutionary change. Just as evolutionary models explain social patterns from the bottom up, learning models theorize social evolution from the bottom up, based on a cognitive “population” of strategies competing for the attention of each individual.

In general form, game-theoretic learning models consist of a probabilistic decision rule and a learning algorithm in which game payoffs are evaluated relative to an aspiration level and the corresponding choice propensities are updated accordingly. The schematic model is diagrammed in Figure 1.

The first step in Figure 1 is the decision by each player to cooperate or defect. This decision is probabilistic, based on the player’s current propensity to cooperate. The resulting outcome then generates payoffs (R, S, P, or T) that the players evaluate as satisfactory or unsatisfactory relative to their aspiration level. Satisfactory payoffs present a positive stimulus (or reward), and unsatisfactory payoffs present a negative stimulus (or punishment). These rewards and punishments then modify the probability of repeating the associated choice, such that satisfactory choices become more likely to be repeated, whereas unsatisfactory choices become less likely.

Erev, Roth, and others (Roth and Erev 1995; Erev and Roth 1998; Erev, Bereby-Meyer, and Roth 1999) used a learning model to estimate globally applicable parameters from data collected across a variety of human subject experiments. They concluded that “low-rationality” models of reinforcement learning may often provide a more accurate prediction than orthodox game theoretical analysis. Macy (1995) and Flache (1996) have also tested a simple reinforcement learning algorithm in controlled social dilemma experiments and found supporting evidence.

Learning models have also been applied outside the laboratory to mixed-motive games played in everyday life in which cooperation is largely unthinking and automatic, based on heuristics, habits, routines, or norms, such as the propensity to loan a tool to a neighbor, tell the truth, or trouble oneself to vote. For example, Kanazawa (2000) used voting data to show how a simple learning model suggested by Macy (1991) solves the “voter paradox” in public choice theory.

Cognitive game theory also applies to conflict in arenas where forward-looking rationality seems most plausible, such as business transactions and international politics. Fischer and Suleiman (1997) developed a formal voting model that closely resembles reinforcement learning and applied the model to intergroup competitions such as the Israeli-Palestinian conflict. They formalized group interaction as a repeated PD game played by the groups' representatives. They assumed that representatives' chances for reelection are influenced by the payoffs that their decisions generate for group members. Successful aggression (*DC*) and peace (*CC*) are evaluated as positive outcomes, whereas mutual conflict (*DD*) and defeat (*CD*) are regarded as unsatisfactory. The negative payoffs from mutual conflict cause representatives from both sides to lose electoral support. If both sides come to favor peace, a period of stability will ensue. However, if only one side favors peace and the other exploits this "weakness," the relationship will move back into mutual conflict. This in turn will eventually undermine support for the leaders of both sides, regenerating peace movements that can become self-reinforcing, but only if both sides avoid the temptation to exploit an opponent who blinks first.

Homans (1974) used William the Conqueror's decision to invade England in 1066 to show how learning theory can explain strategic decisions that might otherwise seem puzzling. From a forward-looking rational choice perspective, William should have refrained from invasion given the enormous uncertainties such as "the danger of a sea voyage, of landing on a hostile shore, and of battle with an English army under the experienced and hitherto successful command of Harold Godwinsson" (p. 46). However, William looked back on a near-perfect, twenty-year record of military victories that reinforced his confidence in his ability to succeed, whatever the odds.

Although interest in cognitive game theory is clearly growing, recent studies have been divided by disciplinary boundaries between sociology, psychology, and economics that have obstructed the theoretical integration needed to isolate and identify the learning principles that underlie new solution concepts. Our study aims to move toward such an integration. We align two prominent models in cognitive game theory and integrate them into a general model of adaptive behavior in two-person, mixed-motive games (or social dilemmas). In doing so, we will identify a set of model-independent learning principles that are necessary and sufficient to generate cooperative solutions. We begin by briefly summarizing the basic principles of reinforcement learning and then introduce two formal specifications of these principles: the Bush-Mosteller stochastic learning model and the Roth-Erev payoff matching model.

LEARNING THEORY AND THE LAW OF EFFECT

Learning theory bridges game theory and biology with the assumption that cooperative decisions are shaped by egoistic reinforcements. Like evolutionary selection, reinforcement is backward-looking (Macy 1990) in that the search for solutions is driven by experience rather than the forward-looking calculation assumed in analytical game theory (Fudenberg and Levine 1998; Weibull 1998). Thorndike (1898) formulated this as the law of effect based on the cognitive psychology of William James. If a behavioral response has a favorable outcome, the neural pathways that triggered

the behavior are strengthened, which “loads the dice in favor of those of its performances which make for the most permanent interests of the brain’s owner” (James [1890] 1981, 143). This connectionist theory anticipates the error back-propagation used in contemporary neural networks (Rummelhart and McClelland 1988). These models show how highly complex behavioral responses can be acquired through repeated exposure to a problem. For example, with a little genetic predisposition and a lot of practice, we can learn to catch a ball while running at full speed without having to stop and calculate the trajectory (as the ball passes overhead).

More precisely, learning theory relaxes three key behavioral assumptions in analytical game-theoretic models of decision:

- Proximity replaces causality as the link between choices and payoffs.
- Reward and punishment replace utility as the motivation for choice.
- Melioration replaces optimization as the basis for the distribution of choices over time.

Proximity, not causality. Compared to analytical game theory, the law of effect imposes a lighter cognitive load on decision makers by assuming experiential induction rather than logical deduction. Players explore the likely consequences of alternative choices and develop preferences for those associated with better outcomes, even though the association may be coincident, superstitious, or causally spurious.

Reward and punishment, not utility. Learning theory differs from game-theoretic utility theory in positing two distinct cognitive mechanisms that guide decisions toward better outcomes, *approach* (driven by reward) and *avoidance* (driven by punishment). The distinction means that aspiration levels are very important for learning theory. The effect of an outcome depends decisively on whether it is coded as gain or loss, satisfactory or unsatisfactory, pleasant or aversive.

Melioration, not optimization. Melioration refers to suboptimal gradient climbing when confronted with “distributed choice” (Herrnstein and Drazin 1991) across recurrent decisions. A good example of distributed choice is the decision whether to cooperate in an iterated social dilemma game. Suppose each side is satisfied when the partner cooperates and dissatisfied when the partner defects. Melioration implies a tendency to repeat choices with satisfactory outcomes even if other choices have higher utility, a behavioral tendency March and Simon (1958) call “satisficing.” In contrast, unsatisfactory outcomes induce a search for alternative outcomes, including a tendency to revisit alternative choices whose outcomes are even worse, a pattern we call “dissatisficing.” Although satisficing is suboptimal when judged by conventional game-theoretic criteria, it may be more effective in leading actors out of social traps than if they were to use more sophisticated decision rules, such as “testing the waters” to see if they could occasionally get away with cheating.

Nevertheless, the law of effect does not solve the social dilemma, it merely reframes it: Where the penalty for cooperation is larger than the reward, and the reward for aggres-

sive behavior is larger than the penalty, how can penalty-averse, reward-seeking actors elude the trap of mutual punishment?

THE BUSH-MOSTELLER STOCHASTIC LEARNING MODEL

The earliest answer was given by Rapoport and Chammah (1965), who used learning theory to propose a Markov model of PD with state transition probabilities given by the payoffs for each state based on the assumption that each player is satisfied only when the partner cooperates. With choice probabilities updated after each move based on the law of effect, mutual cooperation is an absorbing state in the Markov chain.

Macy (1990, 1991) elaborated Rapoport and Chammah's (1965) analysis using computer simulations of their Bush-Mosteller stochastic learning model. Macy identified two learning-theoretic equilibria in the PD game, corresponding to each of the two learning mechanisms, approach and avoidance. Approach implies a self-reinforcing equilibrium (SRE) characterized by satisficing behavior. The SRE obtains when a strategy pair yields payoffs that are mutually rewarding. The SRE can obtain even when both players receive less than their optimal payoff (such as the R payoff for mutual cooperation in the PD game), as long as this payoff exceeds aspirations. Avoidance implies an aversive self-correcting equilibrium (SCE) characterized by dissatisficing behavior. Dissatisficing means that both players will try to avoid an outcome that is better than their worst possible payoff (such as P in the PD game), as long as this payoff is below aspirations. The SCE obtains when the expected change of probabilities is zero and there is a positive probability of punishment as well as reward. This happens when outcomes that reward cooperation or punish defection (causing the probability of cooperation to increase) balance outcomes that punish cooperation or reward defection (causing the probability to decline).³ At equilibrium, the dynamics pushing the probability higher are balanced by the dynamics pushing in the other direction, like a tug-of-war between two equally strong teams.

These learning theoretic equilibria differ fundamentally from the Nash predictions in that agents have an incentive to unilaterally change strategy. In the SRE, both players have an incentive to deviate from unconditional mutual cooperation, but they stay the course as long as the associated payoff exceeds their aspirations. Unlike a pure strategy Nash equilibrium, the SCE has players constantly changing course, but their efforts are self-defeating. It is not that everyone decides he is doing the best he can, it is that his efforts to do better set into motion a dynamic that pulls the rug out from under everyone.

Suppose two players in PD are each satisfied only when the partner cooperates, and each starts out with 0 probability of cooperation. They are both certain to defect, which then causes both probabilities to increase (as an avoidance response). Paradoxically, an increased probability of cooperation now makes a unilateral outcome (CD or DC) more likely than before, and these outcomes punish the cooperator and reward the defector, causing both probabilities to drop. Nevertheless, there is always the chance

3. The expected change in the probabilities of cooperation can also be zero when all outcomes are rewarding, but this equilibrium is not self-correcting. Rather, it is an unstable saddle point from which any deviation will cause probabilities to move toward one of the self-reinforcing pure strategy equilibria of the game.

that both players will defect anyway, causing probabilities to rise further still. Once both probabilities exceed .5, further increases become self-reinforcing by increasing the chances for another bilateral move, and this move is now more likely to be mutual cooperation instead of mutual defection. In short, the players can escape the social trap through *stochastic collusion*, characterized by a chance sequence of bilateral moves in a “drunkard’s walk.” A fortuitous string of bilateral outcomes can increase cooperative probabilities to the point that cooperation becomes self-sustaining—the drunkard wanders out of the gully.

THE ROTH-EREV PAYOFF-MATCHING MODEL

More recently, Roth and Erev (Roth and Erev 1995; Erev and Roth 1998; Erev, Bereby-Meyer, and Roth 1999) have proposed a learning-theoretic alternative to the earlier Bush-Mosteller formulation. Their model draws on the “matching law,” which holds that adaptive actors will choose between alternatives in a ratio that matches the ratio of reward. Applied to social dilemmas, the matching law predicts that players will learn to cooperate to the extent that the payoff for cooperation exceeds that for defection, which is possible only if both players happen to cooperate and defect at the same time (given $R > P$). As with Bush-Mosteller, the path to cooperation is a sequence of bilateral moves.

Like the Bush-Mosteller (BM) stochastic learning model, the Roth-Erev (RE) payoff-matching model implements the three basic principles that distinguish learning from utility theory—experiential induction (vs. logical deduction), reward and punishment (vs. utility), and melioration (vs. optimization). The similarity in substantive assumptions makes it tempting to assume that the two models are mathematically equivalent or, if not, that they nevertheless give equivalent solutions.

On closer inspection, however, we find important differences. Each specification implements reinforcement learning in different ways and with different results. Without a systematic theoretical alignment and integration of the two algorithms, it is not clear whether and under what conditions the backward-looking solution for social dilemmas identified with the BM specification generalizes to the RE model.

Those assumptions can be brought to the surface by close comparison of competing models and by integrating alternative specifications as special cases of a more general model. This is especially important for models that must rely on computational rather than mathematical methods. “Without such a process of close comparison, computational modeling will never provide the clear sense of ‘domain of validity’ that typically can be obtained for mathematized theories” (Axtell et al. 1996, 123).

Unfortunately, learning models have been divided by disciplinary boundaries between sociology, psychology, and economics that have obstructed the theoretical integration needed to isolate and identify the learning principles that underlie new solution concepts. Accordingly, this study aims to align two prominent models in cognitive game theory and integrate them into a general model of adaptive behavior in two-person, mixed-motive games. By “docking” (Axtell et al. 1996) the BM stochastic learning model with the RE payoff-matching model, we can identify learning principles that generalize beyond particular specifications and, at the same time, uncover

hidden assumptions that explain differences in the outcomes and constrain the generality of the theoretical derivations.

In the following section, we formally align the two learning models and integrate them into a general reinforcement learning (GRL) model, for which the BM and RE models are special cases. This process brings to the surface a key hidden assumption with important implications for the generality of learning-theoretic solution concepts. The third section then uses computer simulations of the integrated model to explore the effects of model differences across a range of two-person social dilemma games. The analyses confirm the generality of stochastic collusion but also show that its determinants depend decisively on whether BM or RE assumptions about learning dynamics are employed.

THE BUSH-MOSTELLER AND ROTH-EREV SPECIFICATIONS

In general form, both the BM and RE models implement the stochastic decision process diagrammed in Figure 1, in which choice propensities are updated by the associated outcomes. Thus, both models imply the existence of some aspiration level, A , relative to which cardinal payoffs can be positively or negatively evaluated. Formally, the stimulus s associated with action a is calculated as

$$s_a = \frac{\pi_a - A}{\sup[|T - A|, |R - A|, |P - A|, |S - A|]}, \quad a \in \{C, D\}, \quad (1)$$

where π_a is the payoff associated with action a (R or S if $a = C$, and T or P if $a = D$) and s_a is a positive or negative stimulus derived from π_a . The denominator in equation 1 represents the upper value of the set of possible differences between payoff and aspiration. With this scaling factor, stimulus s is always equal to or less than unity in absolute value, regardless of the magnitude of the corresponding payoff.⁴

Neither model imposes constraints on the determinants of aspirations. Whether aspirations are high or low or change with experience depends on assumptions that are exogenous to both models.

ALIGNING THE TWO MODELS

The two models diverge at the point that these evaluations are used to update choice probabilities. The BM stochastic learning algorithm updates probabilities following an action a (cooperation or defection) as follows:

$$p_{a,t+1} = \begin{cases} p_{a,t} + (1 - p_{a,t})ls_{a,t}, & s_{a,t} \geq 0 \\ p_{a,t} + p_{a,t}ls_{a,t}, & s_{a,t} < 0 \end{cases}, \quad a \in \{C, D\}. \quad (2)$$

4. The Roth-Erev model also does not require that outcomes be normed to $|s| \leq 1$, but norming has no effect and simplifies comparison and integration of the two models.

In equation 2, $p_{a,t}$ is the probability of action a at time t and $s_{a,t}$ is the positive or negative stimulus given by equation 1. The change in the probability for the action not taken, b , obtains from the constraint that probabilities always sum to 1, that is, $p_{b,t+1} = 1 - p_{a,t+1}$. The parameter l is a constant ($0 < l < 1$) that scales the learning rate. With $l \approx 0$, learning is very slow; and with $l \approx 1$, the model approximates a “win-stay, lose-shift” strategy (Catania 1992).

For any value of l , equation 2 implies a decreasing effect of reward as the associated propensity approaches unity but an increasing effect of punishment. Similarly, as the propensity approaches zero, there is a decreasing effect of punishment and a growing effect of reward. This constrains probabilities to approach asymptotically their natural limits.

Like the BM, the RE model is stochastic, but the probabilities are not equivalent to propensities. Propensities are a function of cumulative satisfaction and dissatisfaction with the associated choices, and probabilities are a function of the ratio of propensities. More precisely, the propensity q for action a at time T is the sum of all stimuli s_a a player has ever received when playing a :

$$q_{a,T} = \sum_{t=1}^T s_{a,t}, \quad a \in \{C, D\}. \quad (3)$$

Roth and Erev then use a “probabilistic choice rule” to translate propensities into behavior. The probability p_a of action a at time $t + 1$ is the propensity for a divided by the sum of the propensities at time t :

$$p_{a,t+1} = \frac{q_{a,t}}{q_{a,t} + q_{b,t}}, \quad (a, b) \in \{C, D\}, a \neq b, \quad (4)$$

where a and b represent the binary choices to cooperate or defect. Following action a , the associated propensity q_a increases if the payoff is positive relative to aspirations (by increasing the numerator in equation 4) and decreases if negative. The propensity for b remains constant, but the probability of b declines (by increasing the denominator in the equivalent expression for $p_{b,t+1}$).

An obvious problem with the specification of equations 3 and 4 (but not equation 2) is the possibility of negative probabilities if punishment dominates reinforcement. In their original model, Erev, Bereby-Meyer, and Roth (1999) circumvent this problem with the ad hoc addition of a “clipping” rule. More precisely, their implementation assumes that the new propensity $q_{a,t+1} = q_{a,t} + s_{a,t}$ except for the case where this sum drops below a very small positive constant v . In that case, the new propensity is reset to v . Equations 5 and 6 represent the RE clipping rule by rewriting equation 3 as

$$q_{a,t+1} = q_{a,t} + r_t s_{a,t}, \quad a \in \{C, D\}, \quad (5)$$

where a is the action taken in round t (either cooperation or defection) and r is a response function of the form

$$r_{RE} = \begin{cases} 1, & q_a + s_a \geq 0 \\ \frac{v - q_a}{s_a}, & q_a + s_a < 0 \end{cases}, \quad (a, b) \in \{C, D\}, \quad (6)$$

where all terms except v are indexed on t . This solution causes the response to reinforcement to become discontinuous as propensities approach the lower bound v . To align the model with BM, we replaced the discontinuous function in equation 6 with a more elegant solution. Equation 7 gives asymptotic lower as well as upper limits to the probabilities:

$$r'_{RE} = \begin{cases} \frac{l}{1-l}, & s_a \geq 0 \\ \frac{q_a l}{q_b + q_a l}, & s_a < 0 \end{cases}, \quad (a, b) \in \{C, D\}, a \neq b, \quad (7)$$

where all terms except l (the constant learning rate) are indexed on t . The parameter l sets the baseline learning rate between 0 and 1 (as in equation 2), rather than leaving l implied by the relative magnitude of the payoffs (as in equation 3). More important, equation 7 aligns RE with BM by eliminating the need for a theoretically arbitrary clipping rule. For $s_a \geq 0$, equation 7 is equivalent to equation 6 except that rewards are multiplied by the constant $l/(1-l)$ that increases exponentially with the learning rate l . The important change is for the case where $s_a < 0$. Now r is decreasing in q_a (with a limit at 0). Hence, the marginal effect of additional punishment for action a approaches 0 as the corresponding propensity approaches 0.

Equation 7 aligns RM with BM, allowing easy identification of the essential difference, which is inscribed in equation 4. Rewards increase the denominator in equation 4, depressing the effect on probabilities of unit changes in the numerator. Punishments have the opposite effect, allowing faster changes in probabilities.

To see the significance of this difference, consider first the special case where learning is based on rewards of different magnitude and no stimuli are aversive. Here the RE model corresponds to Blackburn's power law of practice (as cited in Erev and Roth 1998). Erev and Roth (1998) cite Blackburn to support the version of their model that precludes punishment, such that "learning curves tend to be steep initially and then flatten" (p. 859). In the version with both reward and punishment, their model implies what might be termed a power law of learning, in which responsiveness to stimuli declines exponentially with reward and increases with punishment. We have no knowledge of a power law of learning or of any corresponding learning-theoretic concept. For convenience, we have borrowed "fixation" from Freudian psychology.

TABLE 1
Change in Response to Reward and Punishment Following a Recurrent Stimulus

	<i>Following Repeated Reinforcement of C</i>				<i>Following Repeated Punishment of C</i>			
	<i>Response to Reward of</i>		<i>Response to Punishment of</i>		<i>Response to Reward of</i>		<i>Response to Punishment of</i>	
	C	D	C	D	C	D	C	D
BM	Decreases	Increases	Increases	Decreases	Increases	Decreases	Decreases	Increases
RE	Decreases	Decreases	Decreases	Decreases	Increases	Increases	Increases	Increases

NOTE: C = cooperation; D = defection; BM = Bush-Mosteller stochastic learning model; RE = Roth-Erev payoff-matching model.

Fixation is the principal difference between the BM and RE models. Unlike the BM model, which precludes fixation, the RE model builds this in as a necessary implication of the learning algorithm in equation 4. The two models make convergent predictions about

- the declining marginal impact of repeated reinforcement on the response to yet another reward and
- the increasing marginal impact of repeated punishment on the response to an occasional reward.

However, the two models make competing predictions about the effects of repeated punishment on the response to additional punishment and the effects of repeated reward on the response to an occasional punishment, as summarized in Table 1. BM assumes that the marginal impact of punishment decreases with repetition, whereas punishment has its maximum effect following repeated reward for a given action. The power law of learning implies quite the opposite. Repeated failure arouses attention and restores an interest in learning. Hence, the marginal impact of punishment increases with punishment. Conversely, repeated success leads to complacency and inattention to the consequences of behavior, be they reward or punishment.

Note that fixation as implemented in the RE learning model should not be confused with either satisficing or habituation, two well-known behavioral mechanisms that can also promote the routinization of behavior. Fixation differs from satisficing in both causes and effects. Satisficing is caused by a sequence of reinforcements that are associated with a given behavior, which causes the probability of choosing that behavior to approach unity, thereby precluding the chance to find a better alternative. Fixation is caused by a sequence of reinforcements for *any* behavior. Thus, fixation “fixes” any probability distribution over a set of behaviors, including indifference, as long as all behaviors are rewarded. Satisficing and fixation also differ in their effects. Satisficing inhibits search but has no effect on responsiveness to reinforcement. Fixation does not preclude search but instead inhibits responses to the outcomes of all behaviors that may be explored.

Fixation also differs from habituation in both the causes and effects. Habituation is caused by repeated presentation of a stimulus, be it a reward or a punishment. Fixation is caused only by repeated reinforcement; repeated punishment disrupts fixation and restores learning. The effects also differ. Habituation to reward increases sensitivity to punishment, whereas fixation inhibits responsiveness to both reward and punishment.

INTEGRATING THE TWO MODELS

To suppress fixation, Roth and Erev (1995; see also Erev and Roth 1998) added discontinuous functions for forgetting, which suddenly resets the propensity to some lower value. Like the clipping rule, the forgetting rule is theoretically ad hoc and mathematically inelegant. Equation 8 offers a more elegant solution that permits continuous variation in the level of fixation, parameterized as f :

$$r_G = \begin{cases} \frac{l(q_a + q_b)^{1-f}}{1 - l s_a^{1-f}}, & s_a \geq 0 \\ \frac{q_a l(q_a + q_b)^{1-f}}{q_b + q_a l(-s_a)^{1-f}}, & s_a < 0 \end{cases}, \quad (a, b) \in \{C, D\}, a \neq b. \quad (8)$$

As in equation 7, all parameters except f and l are time indexed. The parameter f represents the level of fixation, $0 \leq f \leq 1$. Equation 8 is identical to equation 7 in the limiting case where $f = 1$, corresponding to the RE model, which hardwires fixation.

But notice what happens when $f = 0$. Equation 8 now reduces to the BM stochastic learning algorithm! (The proof is elaborated in the appendix.) By replacing the discontinuous clipping and forgetting functions used by RE with continuous functions, we arrive at a GRL model, with a smoothed version of RE and the original BM as special cases. Hence, the response function in equation 8 is expressed as r_G , denoting the generality of the model. BM and RE are now fully aligned and integrated.

A GENERAL THEORY OF COOPERATION IN SOCIAL DILEMMAS

The GRL model in equations 1, 4, 5, and 8 includes three parameters—aspiration level (A), learning rate (l), and fixation (f)—that can be manipulated to study each of the four mechanisms that we have identified as elements of a learning-theoretic solution concept for social dilemmas: satisficing, dissatisficing, fixation, and random walk.⁵ The model allows us to independently manipulate satisficing (precluded by

5. Habituation to repeated stimuli also affects the learning dynamics and can be modeled with a parameter that scales the degree to which aspiration levels float toward the mean of recent payoffs (Macy and Flache 2002; Erev, Bereby-Meyer, and Roth al 1999). Macy and Flache (2002) explored the consequences of habituation using the Bush-Mosteller algorithm. In this study, we focus on fixation and leave possible interactions with habituation for future research.

high aspirations), dissatisficing (precluded by low aspirations), fixation (precluded by a fixed learning rate), and random walk (precluded by a low learning rate). We can then systematically explore the solutions that emerge with different parameter combinations over each of the three classic types of social dilemma.

THE BASELINE MODEL: STOCHASTIC COLLUSION IN SOCIAL DILEMMAS

We begin by using the GRL model to replicate findings in previous research with the BM specification. This provides a baseline for comparing the learning dynamics in the RE model, specifically, the effects of fixation in interaction with aspiration levels and learning rates.

Macy and Flache (2002) use the BM model to study learning dynamics in three characteristic types of two-person social dilemma games: PD, Chicken, and Stag Hunt. They identify a socially deficient SCE in all three social dilemmas. We can compute SCE analytically by finding the level of cooperation at which the expected change in the probability of cooperation is 0. The expected change is 0 when for both players, the probability of outcomes that reward cooperation or punish defection, weighted by the absolute value of the associated stimuli, equals the probability of outcomes that punish cooperation or reward defection, weighted likewise. With $A = 2$ and the payoff vector $(4, 3, 1, 0)$, the SCE occurs at $p_c = .37$ in PD and at $p_c = .5$ in Chicken and Stag Hunt. Across all possible payoffs, the equilibrium cooperation rate in PD is always below $p_c = .5$, which is the asymptotic upper bound that the solution approaches as R approaches T and P approaches S simultaneously. The lower bound is $p_c = 0$ as P approaches A . In Chicken, the corresponding upper bound is $p_c = 1$ as R approaches T and S approaches A . The lower bound is $p_c = 0$ as R , S , and P all converge on A (retaining $R > A > S > P$). Only in Stag Hunt is it possible that there is no SCE, if $R - T > A - S$. The lower bound for Stag Hunt is $p_c = 0$ as T approaches R and P approaches A .

Using computer simulation, Macy and Flache (2002) show that it is possible to escape this equilibrium through random walk not only in PD but in all dyadic social dilemma games. However, the viability of the solution critically depends on actors' aspiration levels, that is, the benchmark that distinguishes between satisfactory and unsatisfactory outcomes. A self-reinforcing cooperative equilibrium is possible if and only if both players' aspiration levels are lower than the payoff for mutual cooperation (R). Aspirations above this point necessarily preclude a learning theoretic solution to social dilemmas.

Very low aspirations do not preclude mutual cooperation as an equilibrium but may prevent adaptive actors from finding it. If aspiration levels are below maximin,⁶ then mutual or unilateral defection may also be a self-reinforcing equilibrium, even though these outcomes are socially deficient. Once the two players stumble into one of these outcomes, there is no escape, as long as the outcome is mutually reinforcing. If the aspiration level exceeds maximin and falls below R , there is a unique SRE in which both players receive a reward, namely, mutual cooperation.

6. Maximin is the largest possible payoff players can guarantee themselves in a social dilemma. This is punishment in the (P) Prisoner's Dilemma and Stag Hunt and sucker (S) in Chicken.

Figure 2: Stochastic Collusion in Three Social Dilemma Games ($\pi = [4, 3, 1, 0]$, $A = 2$, $l = 0.5$, $f = 0$, $q_{c,1} = q_{d,1} = 1$)

NOTE: A = aspiration level; l = learning rate; f = fixation; q_c = propensity to cooperate; q_d = propensity to defect; R = reward; S = sucker; P = punishment; T = temptation.

If there is at least one SRE in the game matrix, then escape from the SCE is inevitable. It is only a matter of time until a chance sequence of moves causes players to lock in a mutually reinforcing combination of strategy choices. However, the wait may not be practical in the intermediate term if the learning rate is very low. The lower the learning rate, the larger the number of steps that must be fortuitously coordinated to escape the pull of the SCE. Thus, the odds of attaining lock-in increase with the step size in a random walk (see also Macy 1989, 1991).

To summarize, the BM stochastic learning model identifies a social trap—a socially deficient SCE—that lurks inside every PD and Chicken game and most games of Stag Hunt. The model also identifies a backward-looking solution: mutually reinforcing cooperation. However, this solution obtains only when both players have aspiration levels below R , such that each is satisfied when the partner cooperates. Stochastic collusion via random walk is possible only if aspirations are also above maximin and viable for the intermediate term only if the learning rate is sufficiently high.

We can replicate the BM dynamics with the GRL model by setting parameters to preclude fixation and allow random walk, satisficing, and dissatisficing. More precisely,

- Fixation is precluded by setting $f = 0$.
- Aspirations are fixed midway between maximin and minimax. With the payoffs ordered from the set $(4, 3, 1, 0)$ for each of the three social dilemma payoff inequalities, minimax is always 3, maximin is always 1, and $A = 2$. This aspiration level can be interpreted as the expected payoff when behavioral propensities are uninformed by prior experience ($p_a = .5$) such that all four payoffs are equiprobable. This means that mutual cooperation is a unique SRE in each of the three social dilemma games.
- We set the baseline learning rate l high enough to facilitate random walk into the SRE at CC ($l = 0.5$).

Figure 2 confirms the possibility of stochastic collusion in all three social dilemma games, assuming moderate learning rates, moderate aspirations, and no fixation. The figure charts the change in the probability of cooperation, p_c , for one of two players with statistically identical probabilities.

Figure 2 shows how dissatisficing players wander about in a SCE with a positive probability of cooperation that eventually allows them to escape the social trap by random walk. With $f = 0$, the general model reproduces the characteristic BM pattern of stochastic collusion in all three games, but not with equal probability. Mutual cooperation locks in most readily in Stag Hunt and least readily in PD. To test the robustness of this difference, we simulated 1,000 replications of this experiment and measured the proportion of runs that locked into mutual cooperation within 250 iterations. The results confirm the differences between the games. In PD, the lock-in rate for mutual cooperation was only .66, whereas it was .96 in Chicken and 1.0 in Stag Hunt.

These differences reflect subtle but important interactions between aspiration levels and the type of social dilemma—the relative importance of fear (the problem in Stag Hunt) and greed (the problem in Chicken). The simulations also show that satisficing is equally important, at least in PD and Chicken. In these games, appreciation that the payoff for mutual cooperation is “good enough” motivates players to stay the course despite the temptation to cheat (given $T > R$). Otherwise, mutual cooperation will not be self-reinforcing. In Stag Hunt, satisficing is less needed in the long run because there is no temptation to cheat ($R > T$). However, despite the absence of greed, the simulations reveal that even in Stag Hunt, fear may inhibit stochastic collusion if high aspirations limit satisficing. Simulations confirmed that in the absence of fixation, there is an optimal balance point for the aspiration level between satisficing and dissatisficing and between maximin and R (Macy and Flache 2002).

EFFECTS OF FIXATION ON STOCHASTIC COLLUSION

With the BM dynamics as a baseline, we now systematically explore the effects of fixation in interaction with the learning rate and aspirations. Knowledge of these effects will yield a general theory of the governing dynamics for reinforcement learning in social dilemmas, of which BM and RE are special cases.

Fixation has important implications for the effective learning rate, which in turn influences the coordination complexity of stochastic collusion. In both the BM and revised RE models, the effective learning rate depends on the baseline rate (l) and the magnitude of p . Even if the baseline learning rate is high, the effective rate approaches zero as p asymptotically approaches the natural limits of probability. The RE model adds fixation as an additional determinant of effective learning rates. Fixation implies a tendency for learning to slow down with success, and this can be expected to stabilize the SRE. Even if one player should occasionally cheat, fixation prevents either side from paying much attention. Thus, cooperative propensities remain high, and the rewards that induce fixation are quickly restored.

Fixation also implies that learning speeds up with failure. This should destabilize the SCE, leading to a higher probability of random walk into stochastic collusion, all else being equal. This effect of fixation should be even stronger when aspirations are high. The higher actors' aspirations, the smaller are the rewards and the larger the punishments that players can experience. As aspirations approach the R payoff (from below), punishments may become so strong that “de-fixation” propels actors into stochastic collusion following a single incidence of mutual defection.

The lower actors' aspirations, the larger are the rewards and the smaller the punishments that players can experience. Fixation on reward, then, reduces the effective learning rate, thereby increasing the coordination complexity of random walk into the SRE. The smaller the step size, the longer it takes for random walk out of the SCE.

Simply put, when aspirations are above maximin, the power law of learning implies a higher likelihood of obtaining a cooperative equilibrium in social dilemmas compared to what would be expected in the absence of a tendency toward fixation, all else being equal. However, the opposite is the case when aspirations are low.

We used computational experiments to test these intuitions about how fixation interacts with learning rates and aspirations to alter the governing dynamics in each of three types of social dilemma, using a 2 (learning rates: high/low) $\times 2$ (fixation: high/low) $\times 2$ (aspirations: high/low) factorial design with fixation and aspiration levels nested within learning rates. Based on the results for the BM model (with $f = 0$), we set the learning rate l at a level high enough to facilitate stochastic collusion (0.5) and low enough to preclude stochastic collusion (0.05). We begin with a high learning rate to study the effects of fixation on the attraction and stability of the SRE. We then use a low learning rate to look at the effects on the stability of socially deficient SCE.

Within each learning rate, we manipulated fixation using two levels, $f = 0$, corresponding to the BM assumption; and $f = 0.5$, as an approximation of the RE model with a moderate level of fixation. (For robustness, we also tested $f = 1$ but found no qualitative difference with moderate fixation.) Within each level of fixation, we manipulated the aspiration level below and above $A = 2.0$ to see how fixation affects the optimal aspiration level for random walk into mutual cooperation.

We begin by testing the interaction between fixation and satisficing by setting aspirations close to the lower limit at maximin. With $A = 1.05$ and $l = 0.5$, mutual cooperation is the unique SRE in all three games, and stochastic collusion is possible even without fixation. Low aspirations increase the tendency to satisfice, making it more difficult to escape a socially deficient SCE. With low aspirations, the predominance of reward over punishment means that fixation can be expected to increase the difficulty by further stabilizing the SCE.

This intuition is confirmed in Figure 3 across all three games. The dotted lines show the cooperation rate with fixation ($f = 0.5$), and the solid line shows the rate without fixation. Without fixation, low aspirations delay lock-in on mutual cooperation relative to the moderate aspirations in the baseline condition (see Figure 2), but lock-in is still possible within the first 100 iterations. For reliability, we measured the proportion of runs that locked in mutual cooperation within 250 iterations based on 1,000 replications with $A = 1.05$. We found that in PD and Chicken, the lower aspiration level caused a significant decline in the rate of mutual cooperation relative to the baseline condition, but it did not suppress cooperation entirely. In PD, the cooperation rate declined from .66 in the baseline condition to about .14; and for Chicken, the corresponding reduction was from .96 to .83. Only in Stag Hunt did lower aspirations not affect the rate of mutual cooperation.

As expected, fixation considerably exacerbates the problem caused by low aspirations. Figure 3 shows how with $f = 0.5$, lock-in fails to obtain entirely in PD and Chicken and is significantly delayed in Stag Hunt. Reliability tests reveal a highly sig-

Figure 3: Effect of Low Aspirations on Cooperation Rates, with and without Fixation (dotted and solid) in Three Social Dilemmas ($\pi = [4, 3, 1, 0]$, $A = 1.05$, $l = 0.5$, $f = [0, 0.5]$, $q_{c,1} = q_{d,1} = 1$)

NOTE: A = aspiration level; l = learning rate; f = fixation; q_c = propensity to cooperate; q_d = propensity to defect; R = reward; S = sucker; P = punishment; T = temptation.

Figure 4: Effect of High Aspirations on Cooperation Rates, with and without Fixation (dotted and solid) in Three Social Dilemmas ($\pi = [4, 3, 1, 0]$, $A = 2.95$, $l = 0.5$, $f = [0, 0.5]$, $q_{c,1} = q_{d,1} = 1$)

NOTE: A = aspiration level; l = learning rate; f = fixation; q_c = propensity to cooperate; q_d = propensity to defect; R = reward; S = sucker; P = punishment; T = temptation.

nificant effect of fixation. Based on 1,000 replications with $f = 0.5$, the rate of lock-in within 250 iterations dropped to 0 in all three games.

With high aspirations, fixation should have the opposite effect. Without fixation, high aspirations increase the tendency to dissatisfice (overexplore), which should destabilize the SRE. However, the predominance of punishment over reward means that de-fixation can be expected to promote cooperation by negating the destabilizing effects of dissatisficing.

This intuition is confirmed in Figure 4. For symmetry with Figure 3, we assumed an aspiration level of $A = 2.95$, just below the minimax payoff. Without fixation, stochastic collusion remains possible but becomes much more difficult to attain, due to the absence of satisficing.

Figure 4 reveals an effect of fixation that is similar across the three payoff structures. Without fixation, high aspirations preclude stochastic collusion in PD and Chicken. In Stag Hunt, mutual cooperation remains possible but is delayed relative to the baseline condition in Figure 2. With high aspirations, fixation causes the effective learning rate to increase within the first 20 or so iterations, up to a level that is sufficient to quickly obtain lock-in on mutual cooperation. Reliability tests show that the pattern

Figure 5: Effects of Aspiration Level (horizontal axis) on Stochastic Collusion within 250 Iterations, with and without Fixation (dotted and solid) in Three Social Dilemmas ($\pi = [4, 3, 1, 0]$, $l = 0.5$, $f = [0, 0.5]$, $q_{c,1} = q_{d,1} = 1$)

NOTE: l = learning rate; f = fixation; q_c = propensity to cooperate; q_d = propensity to defect; R = reward; S = sucker; P = punishment; T = temptation.

is highly robust. Based on 1,000 replications without fixation, the rate of stochastic collusion within 250 iterations was 0 in PD and Chicken and .64 in Stag Hunt. With fixation ($f = 0.5$), this rate increased to nearly 1 in all three games (.997, .997, and .932 in PD, chicken, and stag hunt, respectively).

To further test the interaction between aspirations and fixation, we varied the aspiration level A across the entire range of payoffs (from 0 to 4) in steps of 0.1. Figure 5 reports the effect of aspirations on the rate of stochastic collusion within 250 iterations, based on 1,000 replications with (dotted line) and without (solid line) fixation.

Figure 5 clearly demonstrates the interaction effects. In all three games, fixation reduces cooperation at low aspiration levels and increases cooperation at high aspiration levels. Fixation also shifts the optimal balance point for aspirations to a level close to the R payoff of the game, or $A = 3$ in PD and Chicken and $A = 4$ in Stag Hunt. In addition, the figure reveals that the interaction between fixation and low aspirations does not depend on whether aspirations are below or above the maximin payoff. As long as aspirations levels fall below $A = 2$, moderate fixation ($f = 0.5$) suppresses random walk into mutual cooperation.

However, the mechanisms differ depending on whether aspirations are above or below maximin. Below maximin, the predominant self-reinforcing equilibria are deficient outcomes: mutual defection in PD and Stag Hunt and unilateral cooperation in Chicken. These equilibria prevail at low aspiration levels because their coordination complexity is considerably lower than that of mutual cooperation. Accordingly, fixation reduces cooperation in this region because it increases the odds that learning dynamics converge on the SRE that is easiest to coordinate, at the expense of mutual cooperation. With aspiration levels above maximin, fixation no longer induces convergence on a deficient SRE, but it inhibits convergence on the unique SRE of mutual cooperation.

To sum up, the GRL model shows that stochastic collusion is a fundamental solution concept for all social dilemmas. However, we also find that the viability of stochastic collusion depends decisively on the assumptions about fixation. With high

Figure 6: Effects of Fixation (dotted line) on Self-Correcting Equilibrium with Low Learning Rates and Low Aspirations in Three Social Dilemma Games ($\pi = [4, 3, 1, 0]$, $l = 0.05$, $A = 1.05$, $f = [0, 0.5]$, $q_{c,1} = q_{d,1} = 1$)

NOTE: A = aspiration level; l = learning rate; f = fixation; q_c = propensity to cooperate; q_d = propensity to defect; R = reward; S = sucker; P = punishment; T = temptation.

aspirations, fixation makes stochastic collusion more likely, whereas with low aspirations, it makes cooperation more difficult.

SELF-CORRECTING EQUILIBRIUM AND FIXATION

We now use low learning rates to study the effect of fixation on the stability of the socially deficient SCE. The coordination complexity of random walk increases exponentially as the learning rate decreases. In our baseline condition with $A = 2$ and $f = 0$ but a low learning rate of $l = 0.05$, stochastic collusion is effectively precluded in all three social dilemma games, even after 1,000 iterations and with moderate aspiration levels. We confirmed this by measuring mean cooperation in the 1,000th iteration over 1,000 replications. For all three types of social dilemma, the mean was not statistically different from the SCE derived analytically from the payoffs ($p = .366$ for PD and $p = .5$ for Chicken and Stag Hunt).

We want to know if fixation makes it easier to escape and whether this depends on the level of aspirations. To find out, we crossed fixation ($f = 0$ and $f = 0.5$) with aspirations just above maximin ($A = 1.05$) and just below minimax ($A = 2.95$), just as we did in the previous study of stochastic collusion with high learning rates.

When aspirations are low, the predominance of reward over punishment should cause the effective learning rate to decline as players fixate on repeated reward. With the baseline learning rate already near 0, fixation should have little effect, and this is confirmed in Figure 6, based on 1,000 iterations in each of the three games. Reduction in the effective learning rate merely smooths out fluctuations in players' probability of cooperation, p_c , around an equilibrium level of about $p_c = .13$ in PD and $p_c = .5$ in Chicken and Stag Hunt. Reliability tests showed a slight negative effect of fixation on the equilibrium rate of cooperation. Based on 1,000 replications, we found that in PD and Chicken, the rate of stochastic collusion within 1,000 iterations was 0, with or without fixation. In Stag Hunt, stochastic collusion remained possible without fixation (with a rate of .21), but this rate dropped to 0 with $f = 0.5$. In short, without fixation, the

Figure 7: Effects of Fixation (dotted line) on Self-Correcting Equilibrium with Low Learning Rates and High Aspirations in Three Social Dilemma Games ($\pi = [4, 3, 1, 0]$, $l = 0.05$, $A = 2.95$, $f = [0, 0.5]$ $q_{c,1} = q_{d,1} = 1$)

NOTE: A = aspiration level; l = learning rate; f = fixation; q_c = propensity to cooperate; q_d = propensity to defect; R = reward; S = sucker; P = punishment; T = temptation.

low learning rate of $l = 0.05$ allows players to escape the SCE only in Stag Hunt. Fixation suppresses even this possibility.

With high aspirations—and a predominance of punishment over reward—fixation should have the opposite effect, causing the effective learning rate to increase. This in turn should make stochastic collusion a more viable solution. To test this possibility, we increased the aspiration level to $A = 2.95$.

Figure 7 confirms the expected cooperative effect of fixation when aspirations are high. The increase in the effective learning rate helps players escape the SCE within about 250 iterations in all three games. Reliability tests show a very powerful effect of fixation. With $f = 0$, the rate of stochastic collusion within 1,000 iterations is 0 in all three games. With moderate fixation ($f = 0.5$), stochastic collusion within 1,000 iterations becomes virtually certain in all three games.

As an additional test, we varied the aspiration level A from 0 to 4 in steps of 0.1, exactly as in Figure 5, only this time with a low baseline learning rate. The results confirm that fixation largely cancels out the effects of a low learning rate. Without fixation ($f = 0$), stochastic collusion is effectively precluded by a low learning rate, regardless of aspiration levels. With fixation, the rate of stochastic collusion over 1,000 iterations was virtually identical to that observed in Figure 5 with a high baseline learning rate.

To sum up, without fixation, low baseline learning rates make it very difficult for backward-looking actors to escape the social trap, regardless of aspiration levels. However, if aspirations are high, fixation increases the effective learning rate to the point that escape becomes as likely as it would be with a moderate baseline learning rate. If aspirations are low, fixation reduces the effective learning rate, but because the rate is already too low to attain stochastic collusion, there is little change in the outcome.

DISCUSSION AND CONCLUSION

Concerns about models of cultural adaptation as analogs of genetic selection have led cognitive game theorists to explore learning-theoretic specifications. Two prominent examples are the BM stochastic learning model and the RE payoff-matching

model. Both models identify two new solution concepts for the problem of cooperation in social dilemmas: a socially deficient SCE (or social trap) and an SRE that is usually (but not always) socially efficient. The models also identify the mechanism by which players can escape the social trap—stochastic collusion, based on a random walk in which both players wander far enough out of the SCE that they escape its “gravitational” pull. Random walk, in turn, implies that a principal obstacle to escape is the coordination complexity of stochastic collusion. Thus, these learning models direct attention to conditions that reduce coordination complexity. These include

- temporal Schelling points, such as the holidays that triggered self-organized truces in World War I;
- small-world networks (Watts 1999), which minimize the average number of partners for each player yet still permit locally successful strategies to globally propagate; and
- high learning rates, which reduce the number of steps that must be coordinated in the random walk.

It is here—the effect of learning rates on stochastic collusion—that the two learning models diverge. The divergence might easily go unnoticed, given the theoretical isomorphism of two learning-theoretic models based on the same three fundamental behavioral principles: experiential induction (vs. logical deduction), reward and punishment (vs. utility), and melioration (vs. optimization). Yet each model implements these principles in different ways and with different results. To identify the differences, we aligned and integrated the two models as special cases of a GRL model. The integration and alignment uncovered a key hidden assumption: the power law of learning. This is the curious but plausible tendency for learning to diminish with success and intensify with failure, which we call fixation. Fixation, in turn, impacts the effective learning rate and, through that, the probability of stochastic collusion.

We used computer simulation to explore the effects of fixation on stochastic collusion in three social dilemma games. The analysis shows how the integration of alternative models can uncover underlying principles and lead to a more general theory. Computational experiments confirmed that stochastic collusion generalizes beyond the particular BM specification. However, we also found that stochastic collusion depends decisively on the interplay of fixation with aspiration levels and the baseline learning rate. The GRL model shows that in the absence of fixation, the viability of stochastic collusion is compromised by low baseline learning rates and both low and high aspirations. With low aspirations, actors learn to accept socially deficient outcomes as “good enough.” We found that fixation exacerbates this problem. When rewards dominate punishments, fixation reduces the effective learning rate, thereby increasing the coordination complexity of stochastic collusion via random walk.

With high aspirations, actors may not feel sufficiently rewarded by mutual cooperation to avoid the temptation to defect. Simulations show that fixation removes this obstacle for stochastic collusion as long as aspirations do not exceed the payoff for mutual cooperation. High aspirations cause punishments to dominate rewards. Fixation, then, increases responsiveness to stimuli, facilitating random walk into the basin of attraction of mutual cooperation.

Our exploration of dynamic solutions to social dilemmas is necessarily incomplete. Alignment and integration of the BM and RE models identified fixation as the decisive difference, and we therefore focused on its interaction with aspiration levels to the exclusion of other factors (such as Schelling points and network structures) that also affect the viability of stochastic collusion. We have also limited the analysis to symmetrical, two-person, simultaneous social dilemma games within a narrow range of possible payoffs. Previous work (Macy 1989, 1991) suggests that the coordination complexity of stochastic collusion in PD increases with the number of players and with payoff asymmetry. We leave these complications to future research.

The identification of fixation as a highly consequential hidden assumption also points to the need to test its effects in behavioral experiments, including its interaction with aspiration levels. Both the BM and the RE models have previously been tested experimentally, but these tests did not address conditions that allow discrimination between the models (Macy 1995; Erev and Roth 1998). Our research identifies these conditions. Erev, Roth, and others (Roth and Erev 1995; Erev and Roth 1998; Erev, Bereby-Meyer, and Roth 1999) estimated parameters for their payoff-matching model from experimental data. However, empirical results cannot be fully understood if key assumptions are hidden in the particular specification of the learning algorithm. Their learning algorithm “hardwires” fixation into the model without an explicit parameter to estimate the level of fixation in observed behavior. Theoretical integration of the two models can inform experimental research that tests the existence of fixation and its effects on cooperation in social dilemmas, including predicted interactions with aspiration levels.⁷

Given the theoretical and empirical limitations of this study, we suggest that its primary contribution may be methodological rather than substantive. By aligning and integrating the BM stochastic learning model with the RE payoff-matching model, we identified a hidden assumption—the power law of learning—that has been previously unnoticed despite the prominent position of both models in the game-theoretic literature. Yet the assumption turns out to be highly consequential. This demonstrates the importance of “docking” (Axtell et al. 1996) in theoretical research based on computational models, a practice that remains all too rare. We hope this study will motivate greater appreciation not only of the emerging field of cognitive game theory but also of the importance of docking in the emerging field of agent-based computational modeling.

7. Calibration of payoffs in decision tasks prior to a social dilemma game can be used to manipulate subjects' aspiration levels. By setting reward (R) just above aspirations, we can test whether stochastic collusion obtains as predicted by Roth-Erev (due to the interaction with fixation) but not by Bush-Mosteller.

APPENDIX

We prove that without fixation ($f = 0$), our GRL model is equivalent to the BM stochastic learning model in equations 1 and 2. By extension, we also obtain the original learning mechanism of Roth and Erev as a special case of a BM stochastic learning model with a dynamic learning rate.

To derive the equation for the effective change in choice probabilities, we use the probabilistic choice rule in equation 4 with the propensities that result after the reinforcement learning rule in equation 5 has been applied in round t . Without loss of generality, let a be the action carried out in t . Equation A1 yields the probability for action a in round $t + 1$, $p_{a,t+1}$, as a function of the reinforcement and the propensities in t :

$$p_{a,t+1} = \frac{q_{a,t} + r_t s_{a,t}}{q_{a,t} + q_{b,t} + r_t s_{a,t}}, \quad (a, b) \in \{C, D\}, a \neq b. \quad (\text{A1})$$

Suppose a was rewarded, that is, $s_{a,t} \geq 0$. Substitution of the response function r in equation A1 by r_G with $f = 0$ yields, after some rearrangement, equation A2:

$$p_{a,t+1} = \frac{q_{a,t} + l q_{b,t} s_{a,t}}{q_{a,t} + q_{b,t}}, \quad (a, b) \in \{C, D\}, a \neq b. \quad (\text{A2})$$

In equation A2, we substitute the terms corresponding to the right-hand side of the probabilistic choice rule in equation 4 by the choice probabilities in round t , $p_{a,t}$ and $p_{b,t} = 1 - p_{a,t}$, respectively. Equation A2 then yields the new choice probability as a function of the old choice probability and the reinforcement. Equation A3 shows the rearrangement where $(a, b) \in \{C, D\}$ and $a \neq b$:

$$p_{a,t+1} = \frac{q_{a,t}}{q_{a,t} + q_{b,t}} + l \frac{q_{b,t}}{q_{a,t} + q_{b,t}} s_{a,t} = p_{a,t} + l(1 - p_{a,t}) s_{a,t}. \quad (\text{A3})$$

The right-hand side of equation A3 is equivalent to the BM updating of probabilities following “upward” reinforcement (or increase of propensity) defined in equation 2.

The BM equations for punishments of a are obtained in the same manner. Let $s_{a,t} < 0$. Then substitution of the response function r_G with $f = 0$ in equation A1 yields, after some rearrangement, the new probability for action a given in equation A4:

$$p_{a,t+1} = \frac{q_{a,t}(1 + l s_{a,t})}{q_{a,t} + q_{b,t}}, \quad (a, b) \in \{C, D\}, a \neq b. \quad (\text{A4})$$

Again, substitution of the terms for $p_{a,t}$ and $p_{b,t} = 1 - p_{a,t}$ according to the probabilistic choice rule yields the new probability as the function of the preceding probability and the punishment that is specified by equation A5.

$$p_{a,t+1} = \frac{q_{a,t}}{q_{a,t} + q_{b,t}} + l \frac{q_{b,t}}{q_{a,t} + q_{b,t}} s_{a,t} = p_{a,t} + l p_{a,t} s_{a,t}. \quad (\text{A5})$$

The rightmost expression in equation A5 is identical with the BM stochastic learning rule for “downward” reinforcement (or reduction of propensity) defined in equation 2.

We now show that the original RE learning mechanism with response function r_{RE} can also be obtained as a special case of a BM learning algorithm with a dynamic learning rate. Consider the case where action a was taken in t and rewarded ($s_a \geq 0$). In the BM algorithm, the new probability $p_{a,t+1}$ for action a is given by equation 2. To obtain the corresponding RE learning rule, we express the learning rate in the BM equation as a function of the present propensities and the most recent payoff. More precisely, the BM equation 2 expands to the RE equations 4 and 5 if the learning rate l is replaced as follows:

$$l_t = \frac{r_{RE,t}}{q_{a,t} + q_{b,t} + r_{RE,t} s_{a,t}}, \quad (a, b) \in \{C, D\}, a \neq b. \quad (A6)$$

Equation A6 shows that for rewards, the RE model is identical to the BM, with a learning rate that declines with the sum of the net payoffs for both actions in the past.

To obtain the BM equation for punishment, the corresponding learning rate function for punishment of action a is

$$l_t = \frac{q_{b,t} r_{RE,t}}{q_{a,t} (q_{a,t} + q_{b,t} + r_{RE,t} s_{a,t})}, \quad (a, b) \in \{C, D\}, a \neq b. \quad (A7)$$

Equation A7 modifies the learning rate as in equation A6, but it also ensures that the learning rate is exponentially amplified as the propensity for the action taken approaches zero. This forces the behavior of the original RE model onto the BM equations. Now the corresponding action propensity drops discontinuously down to its lower bound, a behavior that the BM model with a constant learning rate l avoids with the dampening term $l s_a p_a$, as p_a moves toward zero. Conversely, equation A7 decreases the learning rate as the propensity for the alternative action b approaches zero. The latter effect of equation A7 obtains when the probability for action a is close to one. In that case, the BM model with a constant learning rate l implies a rapid decline of the propensity of a after punishment, unlike the RE model. Equation A7 again forces RE behavior onto the BM equations, as the declining learning rate modifies the strong effect of punishment in this condition.

REFERENCES

- Aunger, R. 2001. *Darwinizing culture: The status of memetics as a science*. Oxford: Oxford University Press.
- Axelrod, R. 1984. *The evolution of cooperation*. New York: Basic Books.
- Axelrod, R., and M. Cohen. 2000. *Harnessing complexity: Organizational implications of a scientific frontier*. New York: Basic Books.
- Axtell, R., R. Axelrod, J. Epstein, and M. Cohen. 1996. Aligning simulation models: A case study and results. *Computational and Mathematical Organization Theory* 1:123-41.
- Blackmore, S. 1999. *The meme machine*. Oxford: Oxford University Press.
- Catania, A. C. 1992. *Learning*. 3d ed. Englewood Cliffs, NJ: Prentice Hall.
- Cohen, M. D., R. L. Riolo, and R. Axelrod. 2001. The role of social structure in the maintenance of cooperative regimes. *Rationality and Society* 13:5-32.
- Dawes, R. M. 1980. Social dilemmas. *Annual Review of Psychology* 31:169-93.
- Dawes, R. M., and R.H. Thaler. 1988. Anomalies: Cooperation. *Journal of Economic Perspectives* 2:187-97.
- Dawkins, R. 1989. *The selfish gene*. 2d ed. Oxford: Oxford University Press.

- Dennett, D. 1995. *Darwin's dangerous idea: Evolution and the meanings of life*. New York: Simon & Schuster.
- Erev, I., Y. Bereby-Meyer, and A. E. Roth. 1999. The effect of adding a constant to all payoffs: Experimental investigation and implications for reinforcement learning models. *Journal of Economic Behavior and Organization* 39:111-28.
- Erev, I., and A. E. Roth. 1998. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* 88:848-79.
- Fischer, I., and R. Suleiman. 1997. Election frequency and the emergence of cooperation in a simulated intergroup conflict. *Journal of Conflict Resolution* 41:483-508.
- Flache, A. 1996. *The double edge of networks: An analysis of the effect of informal networks on cooperation in social dilemmas*. Amsterdam: Thesis Publishers.
- Fudenberg, D., and D. Levine. 1998. *The theory of learning in games*. Cambridge, MA: MIT Press.
- Herrnstein, R. J., and P. Drazin. 1991. Meliorization: A theory of distributed choice. *Journal of Economic Perspectives* 5 (3): 137-56.
- Homans, G. C. 1974. *Social behavior: Its elementary forms*. New York: Harcourt Brace Jovanovich.
- James, W. [1890] 1981. *Principles of psychology*. Cambridge, MA: Harvard University Press.
- Kanazawa, S. 2000. A new solution to the collective action problem: The paradox of voter turnout. *American Sociological Review* 65:433-42.
- Liebrand, W. B. G. 1983. A classification of social dilemma games. *Simulation & Games* 14:123-38.
- Macy, M. W. 1989. Walking out of social traps: A stochastic learning model for the Prisoner's Dilemma. *Rationality and Society* 2:197-219.
- . 1990. Learning theory and the logic of critical mass. *American Sociological Review* 55:809-26.
- . 1991. Learning to cooperate: Stochastic and tacit collusion in social exchange. *American Journal of Sociology* 97:808-43.
- . 1995. PAVLOV and the evolution of cooperation: An experimental test. *Social Psychology Quarterly* 58 (2): 74-87.
- Macy, M. W., and A. Flache. 2002. Learning dynamics in social dilemmas. In *Proceedings of the National Academy of Sciences USA* 99:7229-36.
- March, J. G., and H. A. Simon. 1958. *Organizations*. New York: John Wiley.
- Maynard-Smith, J. 1982. *Evolution and the theory of games*. Cambridge: Cambridge University Press.
- Peyton Young, H. 1998. *Individual strategy and social structure. An evolutionary theory of institutions*. Princeton, NJ: Princeton University Press.
- Rapoport, A. 1994. Editorial comments. *Journal of Conflict Resolution* 38:149-51.
- Rapoport, A., and A. M. Chammah. 1965. *Prisoner's Dilemma: A study in conflict and cooperation*. Ann Arbor: University of Michigan Press.
- Raub, W. 1988. Problematic social situations and the large number dilemma. *Journal of Mathematical Sociology* 13:311-57.
- Roth A. E., and I. Erev. 1995. Learning in extensive-form games: Experimental data and simple dynamic models in intermediate term. *Games and Economic Behavior*, Special Issue: Nobel Symposium 8, pp. 164-212.
- Rummelhart, David E., and James L. McClelland. 1988. *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.
- Schlag, K. H., and G. B. Pollock. 1999. Social roles as an effective learning mechanism. *Rationality and Society* 11:371-97.
- Thorndike, E. L. 1898. *Animal intelligence. An experimental study of the associative processes in animals* (Psychological Monographs 2.8). New York: Macmillan. (Reprint, New Brunswick, NJ, Transaction Publishing, 1999; page reference is to original edition)
- Watts, D. J. 1999. *Small worlds: The dynamics of networks between order and randomness*. Princeton, NJ: Princeton University Press.
- Weibull, J. W. 1998. Evolution, rationality and equilibrium in games. *European Economic Review* 42:641-49.