University of Groningen

Reply

Macy, M.W.; Flache, A.

Published in:
 American Journal of Sociology

DOI:
[10.1086/512963](https://doi.org/10.1086/512963)

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:
 2007

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

Citation for published version (APA):

Macy, M. W., & Flache, A. (2007). Reply: Collective Action and the Empirical Content of Stochastic Learning Models. *American Journal of Sociology*, 112(5), 1546-1554. <https://doi.org/10.1086/512963>

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.

REPLY: COLLECTIVE ACTION AND THE EMPIRICAL CONTENT OF STOCHASTIC LEARNING MODELS¹

We are grateful for the opportunity that Bendor, Diermeier, and Ting (hereafter BDT) have provided to address important questions about the empirical content of learning theoretic solutions to the collective action problem. They discuss two well-known classes of adaptive models—stochastic learning models like that of Bush and Mosteller (1955) that have been applied to collective action and social exchange (Macy 1990, 1991*a*, 1991*b*, 1993; Macy and Flache 2002; Flache and Macy 2002) and theories of satisficing in organizational behavior advanced by Simon (1955) and Cyert and March (1963).

According to Popper (1974, p. 986), when ad hoc assumptions are introduced “to explain a particular difficulty” and these assumptions “cannot be tested independently,” the theory is immunized from refutation. A theory that can explain anything can explain nothing, hence it lacks empirical content. Popper pointed to Darwinism as an example of a theory that is “almost tautological” (Popper 1978). Consequentialist explanations—such as those based on rational choice, natural selection, or reinforcement learning—lack empirical content when any outcome can be explained by proposing some preference, fitness, or aspiration that cannot be tested independently of the outcomes they explain.

BDT argue that learning models in particular lack empirical content if any empirically observed stable outcome—such as unilateral cooperation in the prisoner’s dilemma (PD) game—can be accounted for by invoking some unknown aspiration level, whether aspirations are exogenous or endogenous. The solution, the authors suggest, is to allow both propensities and aspirations to adapt to the outcomes, and to “either introduce enough noise so that agents will be shaken out of arbitrary patterns of behavior, or one can keep the models deterministic but make aspirations depend on social comparisons” (in this issue, p. 1543). They show (BDT 2004) that under these conditions, and given sufficient time, learning mod-

¹ The authors gratefully acknowledge support during the preparation of this reply from the NSF (0537606 and 0433086) to the first author and the Innovational Research Incentive Scheme of the Netherlands Organization for Scientific Research to the second author. Direct correspondence to Michael Macy, Department of Sociology, Cornell University, Ithaca, New York 14853. E-mail: mwm14@cornell.edu or a.flache@rug.nl

Permission to reprint this essay may be obtained only from the author.

els can eventually generate a unique probability distribution that does not depend on assumptions about initial aspirations.

We agree with the authors that “theories of adaptive learning are central to the project of providing microfoundations for sociological theories of macrophenomena such as collective action” (p. 10). We also agree that analytical game theory can take us only so far (see Macy and Flache 2002; BDT 2004). These limitations have led to an explosion of interest in backward-looking models, almost entirely focused on evolutionary game theory. In contrast, surprisingly little attention has been paid to a learning theoretic alternative to evolutionary selection. Along with Roth and Erev (1995; Erev and Roth 1998) and Nowak and Sigmund (1993), BDT are among the pioneers in this important area of game theoretical research using learning models, and we appreciate the obvious effort the authors have invested to improve the empirical content. And we recognize the potential for tautology in consequentialist theories of action, including learning models as well as rational choice and evolution. However, a careful examination of their theorems 1 and 3 and their proposed solutions leads us to conclude that their criticism of the class of learning models that includes the Bush-Mosteller model is unwarranted and that the alternatives they propose have no greater empirical content. More precisely, their argument requires three assumptions:

1. *Aspirations cannot be tested independently.* We show that they can.
2. *Predicted outcomes of the Bush-Mosteller model depend on initial aspirations.* We identify predictions that do not.
3. *BDT’s proposed solutions are not vulnerable to ad hoc explanations.* We discovered a hidden assumption in their model on which the predictions decisively depend and which is at least as difficult to test as are assumptions about aspiration levels.

ASPIRATIONS CAN BE INDEPENDENTLY TESTED

In reference to our earlier work (Macy and Flache 2002) BDT acknowledge, “If in an experiment we could induce fixed aspirations in the (P, R) interval (where P is the payoff to mutual defection and R is the payoff for mutual cooperation), then their prediction that mutual cooperation is the only stable outcome [of the PD game] would be testable” (in this issue, p. 1541, n. 12). BDT offer no evidence to support their assumption that aspirations cannot be induced. Instead, they proceed directly to the conclusion that the impossibility to do so “creates a temptation for the analyst to use ad hoc maneuvers (‘Ah ha! So the agent must have had an aspiration level of such-and-such’)” (p. 1541, n. 12). They ignore the large and growing experimental literature on framing that followed in the wake of Kah-

neman and Tversky's (1984) pioneering study. Research on framing focuses directly on the problem of predicting when an outcome will be coded as a gain or loss, and these studies suggest several ways to induce or measure aspirations (e.g., Kahneman et al. 1991). For example, in a PD game, we could inform participants that the goal in this game is to earn as many points as possible, and to do that, you must discover the strategy that gets your partner to cooperate. We can then use standard manipulation checks to make sure the framing message was effective. Another method is to use three *ceteris paribus* PD games in a within-subject design—a "give-some" game (exchange of contributions), a "take-some" game (exchange of withdrawals), and a "give-some, take-some" game. The Bush-Mosteller model predicts faster and more stable lock in on cooperation in the third game and least in the second. Even if we could not manipulate aspirations, it is sufficient if we can measure them, which we can do by asking participants at the beginning and end of the experiment to classify the set of study outcomes as desirable or undesirable. We can then compare the rates of lock in between those with aspirations inside and outside the (P, R) interval.

As an additional check, we can also manipulate the payoffs. BDT's theorem 1 shows that any stable outcome is possible in a learning model with exogenous aspirations, depending on the aspiration level relative to the payoffs (stationary or nonstationary). This is also one of the main points of our 2002 paper (Macy and Flache 2002)—that self-reinforcing equilibria (hereafter SRE) depend on the level of fixed aspirations. However, BDT do not consider the *change* in SRE as payoffs change across treatment conditions. The effect of aspirations on the evaluation of payoffs depends not only on aspirations but also on the payoffs. Thus, we can not only induce and/or measure the aspirations but also manipulate the payoffs that aspirations frame. For example, consider a PD game with payoff T for unilateral defection. Bush-Mosteller predicts a corresponding probability P_{cc} of locking in mutual cooperation within a finite time period. If T is then increased to T' , then *ceteris paribus*, $P'_{cc} \leq P_{cc}$, regardless of the aspiration level. Intuitively, this is because higher rewards and lower punishments for exploiting a cooperator leave the propensity to defect higher than it would otherwise be.²

Even if we can independently test aspirations, one might reasonably ask, would it not be better to use a model that does not require this

² Even if the aspiration level were to increase with T , as an additional test, we can repeat the experiment but this time hold T constant while reducing all other payoffs. More generally, there are enough combinations of ways that multiple payoffs can be manipulated to rule out any remotely plausible rationalization of a null result in every combination by appealing to an aspiration that was "just so."

assumption? We disagree. If we want to know how aspirations affect the ability of adaptive actors to learn to cooperate, it may be more informative to fix aspirations at different locations in the payoff inequality and observe which outcomes are stable, as we did in our 2002 paper (Macy and Flache 2002). One of the main findings was that endogenous aspirations destabilize all SRE, including mutual cooperation. This is the key reason why BDT propose to endogenize aspirations. But our analysis shows that the assumption that aspirations float is far from innocent. The predicted likelihood of stable mutual cooperation in the PD game is lower if we assume that aspirations adapt to experience than if we assume that players are satisfied when the partner cooperates and dissatisfied when the partner defects. More broadly, we identified the range of aspirations within which mutual cooperation is the unique SRE. Below this range, other SRE become possible, including unilateral cooperation and mutual defection. Above this range, no SRE are possible. We agree with BDT that testing these predictions requires empirical knowledge of aspiration levels, but armed with that knowledge, a model that predicts the effects of aspirations on the likelihood of mutual cooperation could have higher empirical content than a model with endogenous aspirations that generates a unique limiting distribution and is therefore silent on the effects of aspirations on the ability of adaptive actors to find their way out of a social trap.

PREDICTIONS THAT DO NOT DEPEND ON INITIAL ASPIRATIONS

In their 2004 paper, BDT prove that a generic learning model will eventually converge on a unique limiting distribution that is independent of initial aspirations and propensities if certain conditions can be met—if aspirations adapt to experience and there is some arbitrarily small amount of noise, such as the stochastic payoffs BDT propose or propensities that are bound away from the limits of probability so that exploration of alternative actions always remains possible. The problem of empirical content arises only if “every player has a propensity of 1,” as BDT assume in their proof of theorem 3. As it turns out, this is not a problem for the Bush-Mosteller model, because the propensity P for an action a increases asymptotically with a positive payoff π_a

$$P_{a,t+1} = P_{a,t} + (1 - P_{a,t})\pi_{a,t}, \quad (1)$$

where π is normed to the unit interval in absolute value. Simply put, propensities approach unity asymptotically, such that a constant reinforcement has a declining effect. From anywhere in the interior of the distribution, propensities approach but do not reach the natural limits of probability within the finite time frame of our computational experiments.

Put differently, so long as initial propensities are bounded away from the limits by any positive epsilon value, they cannot reach those limits. In all our experiments, there remains a positive probability to explore that declines the longer the players remain in the strategy profile, but the probability does not reach zero. In short, the Bush-Mosteller class of learning models does not conform to the assumption in their proofs of theorems 1 and 3 for any finite $t > 0$.

The assumption that behavior can never be entirely determined means that under the Bush-Mosteller learning dynamics, SRE is a limiting state in the sense that it can never be reached from the outside. To actually enter SRE, the players must start there, and once there, they cannot escape. Starting players at SRE only reveals the static properties of the model. To see the dynamics, it is necessary to start them outside. In the limit, SRE has the static property that it is an absorbing state. Approaching the limit, SRE has the dynamic property that the probability to move away declines in the time that players remain in the strategy profile corresponding to the SRE. The dynamic property of SRE is that it is an attractor—the closer you get, the stronger its pull—yet there remains a positive probability to escape within the finite time of a computational experiment. Because SRE cannot be reached, we measured convergence when propensities approached SRE, that is, when the propensities came within a very small epsilon value of the natural limits of probability.

Alternatively, we could have implemented noise as a constant error term, known in game theory as a “trembling hand.” An example is the “win-stay, lose-shift” learning model of Nowak and Sigmund (1993). We could also have implemented noise with stochastic payoffs, as suggested by BDT. Stochastic payoffs and trembling hands have stationary noise, while the Bush-Mosteller model causes noise to decline as the system approaches SRE. Asymptotic noise has an important advantage. The higher the noise level, the harder it becomes to detect the underlying pattern. Noise also reduces the probability, within a given time period, of a random walk into the basin of attraction of the SRE. Bush-Mosteller has the convenient property that we can have initial noise levels sufficient to shake the system out of an arbitrary local attractor yet still be assured that the system will eventually settle down into a metastable state.

In short, our implementation of Bush-Mosteller bounds propensities away from the natural limits of probability by a declining distance that never reaches zero in the finite time of our computational experiments. This assumption that behavior always retains some idiosyncratic component has an important consequence: even if exogenous aspirations are sufficiently low that players are satisfied with every feasible outcome, not every outcome is an attractor in the dynamics of the game. Consider a PD game in which the players initially play *CD* and both players are

satisfied with this outcome. Eventually, the sucker will try D and receive an even higher payoff. This reward increases the likelihood to try D again.³

With endogenous aspirations and a sufficient number of time steps, the Bush-Mosteller model can generate predictions that do not depend on initial aspirations, so long as propensities are not allowed to reach the corners of the distribution. We see this illustrated in our earlier work (Macy and Flache 2002, table 2), which reports statistically identical results in the PD game for initial aspirations between P and R and between P and S , when aspirations are allowed to float. We credit BDT (2004) for having provided an analytical proof that this result generalizes to any initial aspirations and to a larger class of learning models with endogenous aspirations and any arbitrarily small amount of noise.

NO IMPROVEMENT IN EMPIRICAL CONTENT

In their proposed solutions, BDT have replaced one explicit assumption—that aspirations are “such-and-such”—with two new assumptions, one explicit and one hidden, and both more difficult to test independently. The explicit assumption is that aspirations adapt to experience. This assumption is much more difficult to test than the assumption that aspirations are fixed within any particular interval of the payoff inequality. Standard procedures are available to induce a fixed aspiration level, using methods that were developed for research on framing. It is much harder to induce aspirations that float in a manner that is consistent with the axioms required for BDT’s (2004) proof of a unique limiting distribution. It is curious that BDT do not believe fixed aspirations can be induced, but they appear to have no concerns about how to induce aspirations that adapt to experience, which is the axiom on which the proof for a unique limiting distribution decisively depends.

There is another assumption that is needed for any learning model to make predictions that are independent of initial aspirations. BDT never mention it in the comment above or their 2004 paper, but it follows from their reliance on limiting distributions in a Markov chain. It is only in the long run that learning models (Bush-Mosteller included) can approach a unique limiting distribution. It appears that BDT are also aware of this.

³ In contrast, CD can be a long-term stable outcome in BDT’s model of social comparison. Game theory makes no assumption about payoff symmetry. For example, the S_1 payoff to player 1 in a PD game can be higher than the T_2 payoff to player 2. All that is required is that $T_1 > R_1 > P_1 > S_1$, and the same for player 2. Suppose $S_1 > T_2$. Player 1 will then remain satisfied with unilateral cooperation forever. That cannot happen in the Bush-Mosteller model, where player 1 eventually experiments with defection and receives R_1 , which is higher than S_1 . Player 1’s endogenous aspirations then float up, leaving her unsatisfied with being a sucker.

As they earlier noted, "Our stochastic process has a unique limiting distribution. Moreover the process must eventually converge to that unique distribution from any starting point, that is, from any initial configuration of initial propensities and aspirations" (BDT 2004, p. 27). The key word here is "eventually." We are not aware of any empirical methods that can test whether "eventually" has arrived, other than the observation of the outcome that is predicted to eventually occur.

In the short run, even with endogenous aspirations, stochastic payoffs, social comparison, and trembling hands, all predictions for both the BDT and Bush-Mosteller models remain sensitive to initial aspirations. Peyton Young recognized this as a general problem for the application of theories predicting limiting distributions. He noted that "the dynamics of the process in the short run will be strongly influenced by initial conditions." The length of time needed for the long-run distribution to obtain "depends on the size of the stochastic shocks and the degree of correlation between them, the amount of information agents use in making their decisions, and the extent to which they interact in small, close-knit groups. . . . Thus the length of the long run depends crucially on the details of the learning environment . . . it should come as no surprise to find societies that operate for long periods of time in regimes that do not correspond with the long-run predictions of the theory" (Young 2001, p. 146).

Young's warning applies to any stochastic learning model, including not only Bush-Mosteller but also both of the solutions proposed by BDT. To illustrate, we manipulated initial aspirations in a hybrid learning model that combines asymptotic error with propensities that are bound away from the limits by a very small distance (.001). Furthermore, we assumed a low rate of adaptation of aspirations, where in any learning step new aspirations were a weighted mean of previous aspirations (with a weight of 0.999) and of the most recent payoff experienced. This yields a model that is a member of the class of models for which BDT (2004) prove the existence of a unique limiting distribution, a class that also includes models with stochastic payoffs and social comparison. Table 1 shows the proportion of 1,000 independent realizations that locked in on each of the

TABLE 1
SHORT-RUN DEPENDENCE ON INITIAL ASPRIATIONS

	<i>CC</i>	<i>DD</i>	<i>CD/DC</i>	Mean P_{cc}	Mean A
$A_0 = S \dots$	.044	.862	.03	.061	1.09
$A_0 = T \dots$	.407	0	0	.684	2.40

three pure-strategy profiles of the PD game within 5,000 iterations.⁴

Clearly, these results show that in the short term, the outcome strongly depends on whether initial aspirations are at the upper or lower end of the range of feasible payoffs (from T to S). However, when we allowed the experiments to continue for a larger number of time steps, the results approached a unique limiting distribution that is very close to the distribution we observe for $A = T$ in table 1 (proportion $CC = 0.42$; mean $P_{cc} = 0.705$; mean $A = 2.39$). Suppose we were then to observe in an empirical study the results predicted for $A = S$ in table 1, instead of the predicted limiting distribution. Now the ad hoc maneuver is “Ah ha! The number of learning steps in the experiment must have been too small” instead of “an aspiration level that is such-and-such.”

While it is reasonable to assume that exogenous aspirations may be imposed or measured, the assumption that the number of learning steps was sufficient for the empirical test to reach the predicted limiting distribution will be much more difficult to test independently. If we want predictions that do not depend on assumptions about which we have no knowledge, we are far better off with predictions that depend on knowing the aspiration level than with predictions that depend on knowing when “eventually” has arrived.⁵ The problem is complicated further by the dependence not only on initial aspirations and propensities, but also on the form and amount of noise and the rates at which aspirations and propensities adapt. These additional parameters invite still more ad hoc explanations of an inconvenient result.

We close with Popper’s reconsideration of the empirical content of Darwinism: “I have in the past described the theory as ‘almost tautological,’ and I have tried to explain how the theory of natural selection could be untestable (as is a tautology) and yet of great scientific interest. . . . Nevertheless, I have changed my mind about the testability and logical

⁴ We measured lock in to an SRE when propensities to play the SRE strategy were within 0.00001 of the highest possible propensity of 0.999, and the corresponding payoffs exceeded the current aspiration levels.

⁵ With exogenous aspirations, we also do not know how long to wait for lock in to occur. However, exogenous aspirations allow testable predictions about the out-of-equilibrium dynamics, not just the probability distribution of stable outcomes.

status of the theory of natural selection; and I am glad to have an opportunity to make a recantation" (Popper 1978, p. 345). We hope that BDT will follow in Popper's footsteps and reconsider what appears to be a similarly hasty judgment about Bush-Mosteller models of adaptive learning.

MICHAEL MACY

Cornell University

ANDREAS FLACHE

University of Groningen

REFERENCES

- Bendor, J., D. Diermeier, and M. Ting. 2004. "The Empirical Content of Adaptive Models." Research Paper no. 1877. Stanford Graduate School of Business.
- Bush, R. R., and F. Mosteller. 1955. *Stochastic Models for Learning*. New York: Wiley.
- Cyert, Richard, and James March. 1963. *A Behavioral Theory of the Firm*. Englewood Cliffs, N.J.: Prentice Hall.
- 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.
- Flache, Andreas, and Michael Macy. 2002. "Stochastic Collusion and the Power Law of Learning." *Journal of Conflict Resolution* 46:629–53.
- Kahneman, D., J. L. Knetsch, and R. H. Thaler. 1991. "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias." *Journal of Economic Perspectives* 5: 193–206.
- Kahneman, D., and A. Tversky. 1984. "Choices, Values and Frames." *American Psychologist* 39:341–50.
- Macy, Michael. 1990. "Learning Theory and the Logic of Critical Mass." *American Sociological Review* 55:809–26.
- . 1991a. "Chains of Cooperation: Threshold Effects in Collective Action." *American Sociological Review* 56:730–47.
- . 1991b. "Learning to Cooperate: Stochastic and Tacit Collusion in Social Exchange." *American Journal of Sociology* 97:808–43.
- . 1993. "Backward-Looking Social Control." *American Sociological Review* 58: 819–36.
- Macy, Michael, and Andreas Flache. 2002. "Learning Dynamics in Social Dilemmas." *Proceedings of the National Academy of Sciences* 99:7229–36.
- Nowak, M., and K. Sigmund. 1993. "A Strategy of Win-Stay, Lose-Shift that Outperforms Tit-for-Tat in the Prisoner's Dilemma Game." *Nature* 364:56–58.
- Popper, Karl. 1974. "Replies to My Critics." In *The Philosophy of Karl Popper*, 2 vols. Edited by P. A. Schilpp. LaSalle, Ill.: Open Court Press.
- . 1978. "Natural Selection and the Emergence of Mind." *Dialectica* 32:339–55.
- Roth, A. E., and I. Erev. 1995. "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in Intermediate Term." In "Nobel Symposium." Special issue of *Games and Economic Behavior* 164–212.
- Simon, H. 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics* 69 (1): 99–118.
- Young, P. 2001. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton, N.J.: Princeton University Press.