Imitation and Inequity in Avoiding the Tragedy of the Commons

Terry Van Belle and **David H. Ackley**
Department of Computer Science
University of New Mexico
Albuquerque, NM, USA

Abstract

We present a tragedy of the commons model in which individuals have no *a priori* knowledge of the immediate consequences of their actions: Each agent chooses actions using a simple neural network, which it gradually modifies to more closely imitate those of its wealthier neighbors. For a small commons size, the model leads neither to the tragedy of complete resource exhaustion nor to complete cooperation, but instead to the emergence of polarized 'economic classes' of poor and altruistic agents living amongst rich and greedy ones. The tragedy does emerge with larger commons sizes; we found that adding a degree of enforced local sharing among neighbors staves off tragedy there, and once again the economic stratification emerges. Though simple, the model displays a surprising range of dynamic behaviors at multiple temporal and spatial scales, as two fundamentally conflicting 'ideologies' war for control of agent behavior.

The Tragedy of the Commons

The Tragedy of the Commons (Hardin, 1968), like its sister model, the Prisoner's Dilemma (Axelrod and Dion, 1988), is an abstraction of social interactions in which the long-term common good of a group of agents is in conflict with each agent's short-term individual interests. In the classical example, as long as the village's common pasture is not overgrazed, it can provide ample feed to the villagers' livestock, but each individual livestock owner can gain advantage—if only temporarily—by grazing more animals there.

Because of its importance in everyday life, the tragedy has been extensively studied. Many solutions to the problem have been proposed, using both formal models (Axtell et al., 2000; Axelrod, 1997; Riolo et al., 2001; Axtell, 2003), and real-world examples (Milinski et al., 2002; Ostrom, 1990).

It is typically presumed that each agent in a commons model is inherently aware of the relationship between its action and its short-term marginal return—that putting another cow in the common pasture will increase that cow owner's share—and the model assumes that each agent makes (fully or boundedly) rational decisions on that basis. By contrast, we consider a commons model in which each agent chooses its actions based on a simple three weight neural network that maps from agent wealth and local resources to a decision about what fraction of the available resource to consume. Depending on the weights the neural network can produce a significant range of behavioral strategies; we investigate what happens when agents begin with random neural network weights, but then over time modify themselves by imitating the neural networks of the richest people in their local neighborhoods. For some parameter values, we observe neither the tragedy outcome of universal defection nor the 'utopian' outcome of unanimous cooperation, but instead the emergence of polarized economic classes and the exploitation of the obliging poor by the greedy rich.

In such cases, a quasi-stable balance is achieved by an extreme disparity in resource usage, with some in the population maximizing their take, and the remainder minimizing their take to compensate. For fully independent and rational agents, such a situation is unstable because the exploited can increase their resources by being just as greedy as their exploiters. However, in a system in which agents (willingly or under coercion) adopt the behavioral strategies employed by the rich, a system of economic differentiation and exploitation can emerge and spread. We find that in a variety of circumstances in which the tragedy of the commons is avoided, such endemic social inequities often arise and persist.

The Model

The world consists of a two-dimensional 64×64 grid. Each grid square permanently contains exactly one agent. The neighborhood of an agent consists of the 8 immediate squares, or fewer if the agent is on an edge of the world.

Each square initially contains one unit of some spontaneously growing *resource* that is valuable to the agents, but the squares are grouped into many disjoint rectangular *commons*, and the resources are pooled within each commons. For example, if the commons are 2 squares high and 1 square wide, then the 4,096 squares of the grid would be grouped into 2,048 commons, each which contains 2 villagers who have 2 units of resource initially at their disposal. The neighborhood of an agent, however, is *not* restricted by the boundary of the commons, allowing information to flow between commons via imitation. See Figure 1.

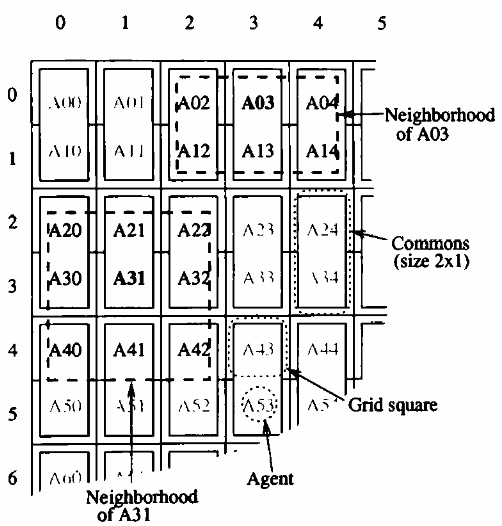

Figure 1: The upper-left corner of a world with 2 × 1 commons. The neighborhoods of agents A_{31} and A_{03} are shown with dashed lines; note that only A_{21} is in the same commons as A_{31}.

The model is run for some specified number of timesteps. During a timestep, the following phases occur effectively simultaneously in all commons: First, the agents may take a certain amount from the currently available resource in their commons. Then agents pay 'taxes', then agents imitate a richer neighbor (if any), then their neural network weights change randomly a bit, and then the commons resource grows.

Take Each agent is controlled by a neural net (Rumelhart and McClelland, 1986) with one output and three inputs: A *bias* input B that is always set to 1, a *resource* input R providing the current average resources per square in the agent's commons, and a *wealth* input W specifying the agent's current holdings in terms of accumulated resource. Each input is scaled by a modifiable link weight (x_B, x_R, x_W), which is initialized to a uniform random value from the range $[-1, +1]$ at the start of a model run. The output of the neuron represents the desired *take* of agent A_i in the current circumstances, and is determined by a sigmoid of the net scaled input:

$$t_i = \frac{1}{1 + e^{-(Bx_B + Rx_R + Wx_W)}}$$

What agent A_i actually gets, g_i, is determined by its desired take t_i, and the current amount of resource R in the commons:

$$g_i = \frac{t_i}{n} R$$

where n is the number of agents sharing the commons. R is then reduced by the sum of the g_i's.

This commons model is inherently 'kinder' than it might be, in the sense that an agent can never take more than a $\frac{1}{n}$th share of the available resources on any timestep. The essential tragic opportunity remains, however, because if everyone in a commons takes their full share, then the resource will be fully exhausted on that turn. Only if at least some agents choose to exercise self-restraint, and take less than their full share, will significant resources be left over to grow in the future.

Tax Next, each agent's wealth W_i is reduced by a combination of a fixed term (parameter τ_{fixed}) to cover the cost of existing, and a proportional 'spoilage' (parameter τ_{spoil}) term:

$$W_i' = (1 - \tau_{spoil}) \cdot \max(W_i - \tau_{fixed}, 0)$$

Imitate Next, each agent A_i considers imitating one of its neighbors and computes an *imitation weight change* (vector $\Delta\mathbf{x}_i$). If all the neighbors A_j have the same or less wealth ($\forall j\ W_i > W_j$), then $\Delta\mathbf{x}_i = \mathbf{0}$. Otherwise, A_i selects randomly from set K of richer neighbors based on how much richer ($\Delta W_j = W_j - W_i$) they are, picking neighbor N_j with probability

$$prob(N_j) = \frac{\Delta W_j}{\sum_{k \in K} \Delta W_k}$$

Given the selected neighbor A_j, the imitation weight change is:

$$\Delta \mathbf{x}_i = \alpha(\mathbf{x}_j - \mathbf{x}_i)$$

with imitation rate parameter α.

Noise Whether or not imitation occurs, a noise vector $N(\mathbf{0}, \beta)$ is drawn from a random Gaussian vector distribution, with mean $\mathbf{0}$ and standard deviation specified by parameter β.

Update Once all agents have performed the previous phases, then finally each agent A_i's weights are updated, based on the (possibly zero) $\Delta \mathbf{x}_i$ and the Gaussian noise:

$$\mathbf{x}'_i = \mathbf{x}_i + \Delta \mathbf{x}_i + N(\mathbf{0}, \beta)$$

Grow In the final timestep phase, in each commons the pooled resource R grows according to a logistic equation, increasing exponentially when R is small, with growth tapering off as the resource per square approaches maximum resource parameter R_{max}. In addition, minimum resource parameter $R_{min} > 0$ gives each commons some chance to recover from poor management:

$$R' = \max(nR_{min},\ R + \gamma R(1 - \frac{R}{nR_{max}}))$$

Experimental Results

Although this is a new model, and many experiments remain to be performed, we have already gained substantial experience exploring the model. Here, first, we discuss results in the degenerate case of a size 1×1 'commons', then focus primarily on data and observations from the smallest non-trivial commons, of size 2×1, and finally touch on results from larger commons. Except for the commons size and timestep limits, model parameter values were held constant for all experiments reported in this paper, and are given in Table 1.

1×1 Commons: Private Ownership

To calibrate our understanding of the model, we tested its behavior with the world divided into 1×1 'commons', so the results of each agent's choices directly affect only itself. Though tapping separate resources, the agents' behaviors are still coupled via the imitation process. Any time an agent overuses its resource, its wealth drops and it will start to

Parameter	Value	Parameter	Value
World width	64	τ_{fixed}	0.1
World height	64	τ_{spoil}	0.1
R_{min}	0.01	α	0.1
R_{max}	10	β	0.1
γ	1.25		

Table 1: Experimental parameter values. See text.

imitate its more temperate neighbors; a similar reaction occurs if it underuses its resource. Without any resource sharing, there is no possibility of cooperation or defection.

Figure 2 displays typical behavior of this model: The average resource per square quickly climbs to around 5 units—just about half of the $R_{max} = 10$ saturation value. The average take settles down to about 0.3, and the average wealth slowly rises past 15. In contrast to this seemingly sedate high-level view of the model evolution, the top of Figure 3 shows a close-up view of the 'take' of a sample agent within the grid. With 1×1 commons, agents develop a 'farming' strategy that involves choosing a very low take for one or a few timesteps, depleting their wealth in the process, until the resource reaches a threshold size, at which point they 'harvest' the sizable bounty, and the cycle begins again.

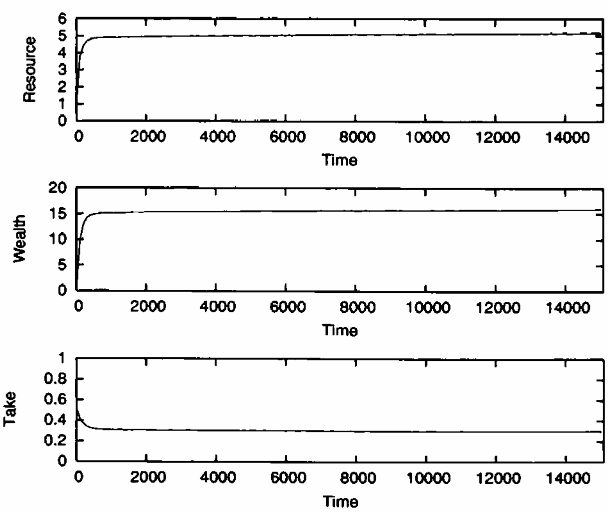

Figure 2: Average common resource levels, agent wealth and take, for 1×1 commons, averaged across 50 runs.

This is not the only sensible strategy in a 1×1 commons, and we have also shown that agents choosing a constant take at an optimal level do significantly better than these emergent farmers,

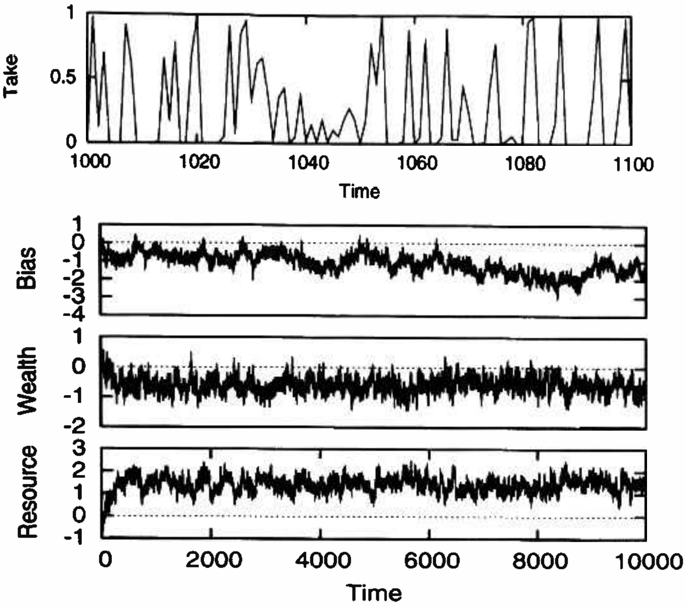

Figure 3: *(Top graph)* A sequence of takes from timesteps 1000–1100, for a sample agent using the 'farming' strategy in a 1 × 1 commons. *(Bottom graphs)* The evolution of the weights of a sample agent in a 1 × 1 commons.

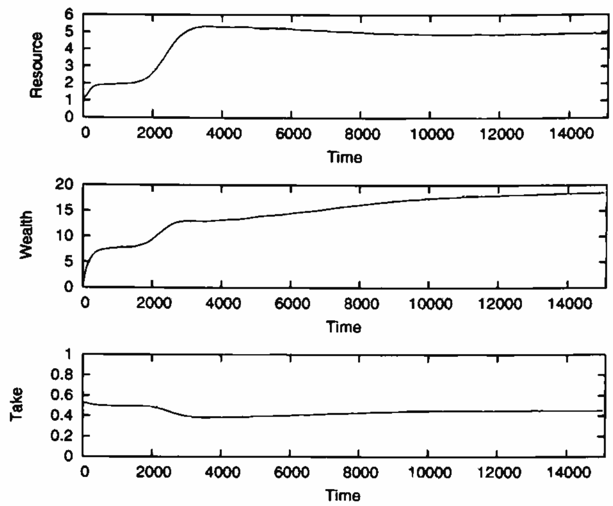

Figure 4: Average common resource levels, agent wealth, and take for 2 × 1 commons, averaged across 50 runs.

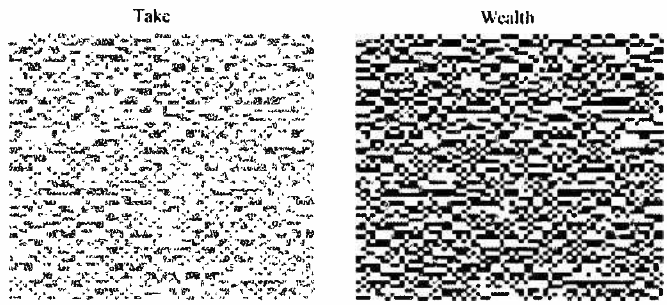

Figure 5: Distributions of t_i (left pane) and W_i (right pane) in a world of 2 × 1 commons. Most commons consist of a rich, greedy agent above or below a poor, altruistic one. See text.

but the farming strategy can be implemented by the neural network controller without requiring precisely chosen weights, as the bottom of Figure 3 suggests: A negative bias tends to shut down the agent take except when the resource multiplied by the resource weight is large enough to overcome it.

2 × 1 Commons: Emergent Inequity

Of course a 1×1 'commons' isn't a commons at all in any significant sense. The 2×1 commons, with two agents sharing each pooled resource, displays far more complex and interesting behaviors. Figure 4 displays overall behavior of this model, averaged over 50 runs varying only the random number seed.[1] The inflection points in Figure 4 suggest temporal phenomena effects at multiple scales, and indeed this seems to be the case.

Consider Figure 5, which presents a snapshot from a 2 × 1 commons run captured at time 50,000, showing the distribution of agent takes (left panel) and wealth (right panel), with darker grays indicating larger values. We see that most of the vertically oriented 2 × 1 commons have polarized so that they contain a single rich, greedy agent and a single poor, altruistic one. Figure 6 reveals the essential

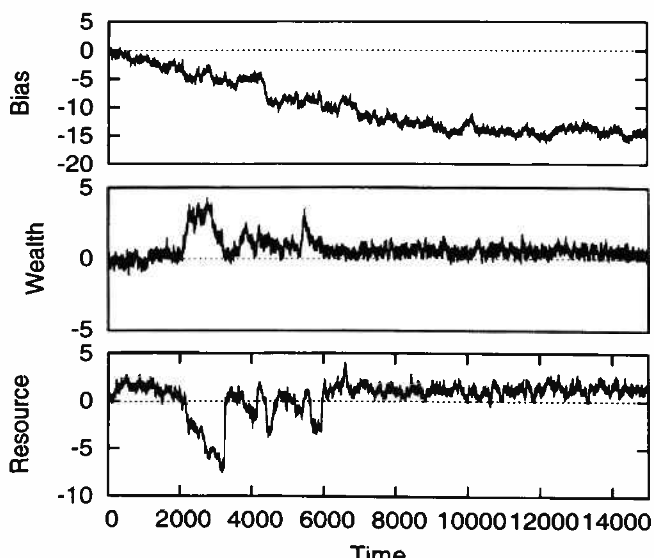

Figure 6: Neural network weights vs. time for a sample agent in a 2 × 1 commons.

[1] A 95% confidence interval around the final values in Figures 2 and 4 is ±0.025 for Resources, ±0.05 for Wealth, and ±0.002 for Take.

trick at work—a negative bias weight encourages taking little, while a positive wealth weight allows agents that are already wealthy to choose a large take. When poor agents imitate rich agents possessing this network, they eventually embrace this "rich get richer" philosophy, even though following it specifically prevents them from realizing the benefits that flow to the rich agent employing the same strategy.

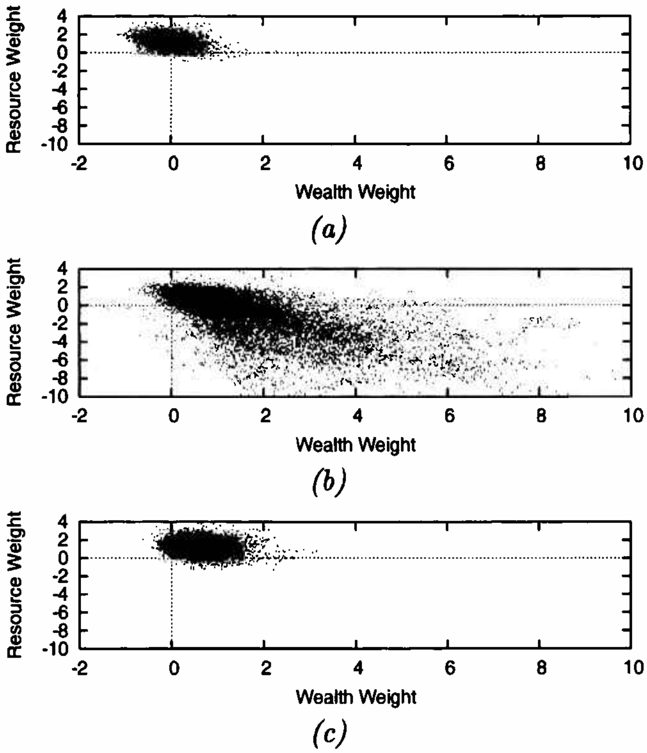

Figure 7: Resource and wealth weights during timesteps 1–2000 *(a)*, 2001–10,000 *(b)*, and 10,001–18,000 *(c)*, for a sample agent in a 2 × 1 commons. Aggregated across 20 runs. Some points in *(b)* are clipped beyond (10, -10).

There is another unexpected aspect of this strategy visible in Figure 6: Particularly in the era around timesteps 2000–6000, the resource weight is frequently actually *negative*, meaning that other things being equal, the more resource is available, the *less* of it you should take. Of course other things aren't equal; for the rich the positive wealth weight is more than enough to overcome this pathological aversion to resources and keep on taking. Over the 20 runs that we have currently examined, it appears this "big lie" strategy eventually gives way to an alternative "rich get richer" strategy employing a large negative bias weight without requiring a negative resource weight. The weight distributions in Figure 7 suggest the progression through three distinct phases is quite robust.

Larger Commons: Avoiding Tragedy With Enforced Sharing

Moving beyond the 2 × 1 commons, we found that similar polarizations can occur in 3 × 1 and 4 × 1 commons, though the class emergence outcome occurs less frequently. In the basic model, with larger or fatter commons such as 2 × 2 and 4 × 4, the usual outcome is the tragedy, ending with all agent takes near 1 while their gets languish near 0.

In an attempt to stave off tragedy in larger commons sizes without moving to a completely centralized approach, we added a neighborhood sharing mechanism to the model: Between the **Take** and **Tax** phases, each agent A_i shares a fraction (parameter η) of its wealth W_i evenly and simultaneously with its n_i neighbors:

$$W_i' = W_i - \eta W_i + \frac{1}{n_i}\sum_{j=1}^{n_i} \eta W_j$$

Setting η to a non-zero value can stabilize the 2 × 2 and 4 × 4 cases, though it takes more sharing to stabilize the larger commons: Sharing as little as 10% ($\eta = 0.1$) can stabilize the 2 × 2 case, while at least $\eta = 0.6$ is necessary for 4 × 4 commons.

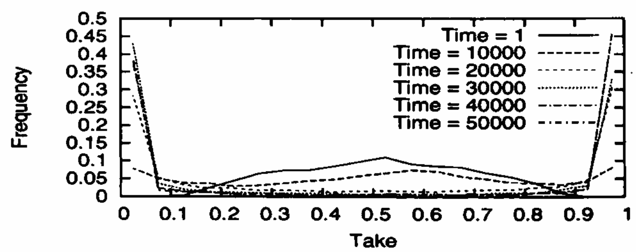

Figure 8: Histogram of the takes for agents in a world with 4 × 4 commons, $\eta = 0.6$.

Again, however, once the tragedy has been averted, the result is social inequity. Figure 8 displays the progressive loss of the middle class in a 4 × 4 commons model, and Figure 9 reveals the characteristic banded patterns associated with exploitation. Each 4 × 4 commons tends to organize itself into roughly half selfish and half altruistic agents, and the selfish agents tend to cluster together along the borders between commons, so that they are close to as many rich, selfish neighbors as possible. This arrangement allows them to share relatively little of their resources with nearby

Figure 9: Distributions of t_i (left pane) and W_i (right pane) in a run with 4×4 resources and $\eta = 0.6$.

poor, and mostly to wash it back and forth between agents that are in a similar situation to themselves.

Discussion and Conclusion

A key aspect of our model is that the imitation process does not focus on the overt behavior of the agent imitated, but on that agent's entire programming. If poor agents imitate only the rich's 'take the max' actions, the tragedy of the commons promptly ensues. Stable patterns of social inequity can emerge because the poorer agents buy into an entire "philosophy," even if that philosophy then dictates they behave differently than those they look to for guidance.

Even when the internal details of an agent's decision-making processes are largely unobservable, machine learning techniques (Rumelhart and McClelland, 1986) may be used for imitation if the inputs associated with the agent's selected actions are available. In a general sense, things like story plots and morality plays are also ways of conveying input-output associations to the receptive; it is easy to speculate on ways that computational dependencies that encourage social inequity may be embedded in some of society's most familiar ideas.

It is important to note that although the exploitation patterns we observed are often quite stable, with an agent remaining rich or poor for many hundreds of consecutive timesteps, such situations are not frozen for all time. Perhaps because of the inescapable **Noise** step in the algorithm, the patterns drift like sand dunes over the course of thousands of timesteps, while maintaining the overall distribution of rich and poor. Sometimes a rich agent *is* brought down and a poor one *does* rise. By contrast, once a model goes firmly for global Hobbesian defection, we have never seen any sustained reemergence of resources and wealth.

The ubiquity of social inequity in these quite simple models suggests that perhaps something fundamental is at work. While it seems a clearly honorable goal to seek methods that avoid the tragedy of the commons with evenhanded fairness and justice, if one is stuck with mechanisms that allow inequities to emerge and persist, perhaps a most immediate goal is to ensure that the inequities remain mobile.

Acknowledgments

The research presented in this paper benefitted greatly from discussion and contributions by the members of the 2003 UNM Artificial Life Research Group. It was supported in part by DARPA contract F30602-00-2-0584, and in part by NSF contract ANI 9986555.

References

[1] Axelrod, R. (1997). *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration.* Princeton University Press, Princeton, NJ.

[2] Axelrod, R. and Dion, D. (1988). The further evolution of cooperation. *Science*, 242:1385–1390.

[3] Axtell, R. (2003). The emergence of institutions of self-governance on the commons. Brookings Institution Working Paper, Center on Social and Economic Dynamics.

[4] Axtell, R., Epstein, J. M., and Young, H. P. (2000). The emergence of classes in a multi-agent bargaining model. Working Paper 9, Center on Social and Economic Dynamics.

[5] Hardin, G. (1968). The tragedy of the commons. *Science*, 162:1243–1248.

[6] Milinski, M., Semmann, D., and Krambeck, H.-J. (2002). Reputation helps solve the 'tragedy of the commons'. *Nature*, 415:424–426.

[7] Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action.* Cambridge University Press, New York.

[8] Riolo, R. L., Cohen, M. D., and Axelrod, R. (2001). Evolution of cooperation without reciprocity. *Nature*, 414:441–443.

[9] Rumelhart, D. E. and McClelland, J. L. (1986). *Parallel Distributed Processing, Volume 1: Foundations.* MIT Press, Cambridge, MA.