A Game-Theoretic and Dynamical-Systems Analysis of Selection Methods in Coevolution

Sevan G. Ficici, Ofer Melnik, and Jordan B. Pollack

***Abstract*—We use evolutionary game theory (EGT) to investigate the dynamics and equilibria of selection methods in coevolutionary algorithms. The canonical selection method used in EGT is equivalent to the standard "fitness-proportional" selection method used in evolutionary algorithms. All attractors of the EGT dynamic are Nash equilibria; we focus on simple symmetric variable-sum games that have *polymorphic* Nash-equilibrium attractors. Against the dynamics of proportional selection, we contrast the behaviors of truncation selection, (μ, λ), $(\mu + \lambda)$, linear ranking, Boltzmann, and tournament selection. Except for Boltzmann selection, each of the methods we test unconditionally fail to achieve polymorphic Nash equilibrium. Instead, we find point attractors that lack game-theoretic justification, cyclic dynamics, or chaos. Boltzmann selection converges onto polymorphic Nash equilibrium only when selection pressure is sufficiently low; otherwise, we obtain attracting limit-cycles or chaos. Coevolutionary algorithms are often used to search for solutions (e.g., Nash equilibria) of games of strategy; our results show that many selection methods are inappropriate for finding polymorphic Nash solutions to variable-sum games. Another application of coevolution is to model other systems; our results emphasize the degree to which the model's behavior is sensitive to implementation details regarding selection—details that we might not otherwise believe to be critical.**

***Index Terms*—Coevolutionary algorithms, dynamical systems, game theory, selection methods, solution concepts.**

I. Introduction

THIS PAPER considers the dynamical and game-theoretic properties of selection processes in coevolutionary algorithms. Like any evolutionary algorithm (EA), a coevolutionary algorithm is a population-based stochastic search method inspired by the principle of natural selection. Unlike an ordinary EA, a *co*evolutionary algorithm evaluates individuals by having them interact with each other, and the outcomes of these interactions determine the individuals' reproductive potentials. The realization of reproductive potential is accomplished by a *selection method*, which maps a set of (possibly nonscalar) evaluation results into a set of nonnegative scalars that specify the reproductive success of each individual, measured in numbers of offspring; this mapping is often stochastic.

Many design choices go into the construction of a coevolutionary algorithm, and each design choice potentially impacts the algorithm's behavior in significant ways, often by interacting nonlinearly with other design choices. Each selection method expresses selection pressure differently. The efficacy of any given selection method (e.g., with respect to genetic diversity, convergence, and overall system performance) depends upon other design choices, such as the genetic representation and variation operators used in a system; this nonlinearity ultimately requires a variety of selection methods to be handy.

Selection methods convert an ordering (possibly a partial ordering) of evaluation results into a total ordering of expected numbers of offspring. While different selection methods use different mappings, the expected number of offspring for some individual A will never be less than the expected number of offspring for some other individual B if evaluation determines that individual A is a superior individual to B. (Coevolution research does include work that penalizes overly successful individuals in zero-sum games, e.g., [1] and [2]; we take an individual's evaluation score to reflect all such additional processing steps.) Thus, on a fundamental level, selection methods are mutually consistent in their operation. This consistency suggests a degree of interchangeability between selection methods (their unique nonlinear interactions with the rest of the system notwithstanding), at least for the simple framework in which we will examine them, if not in real-world algorithms.

Coevolutionary dynamics are notoriously complex. To focus on our attention on selection dynamics, we will use a simple evolutionary game-theoretic framework [3] to eliminate confounding factors such as those related to genetic variation, noisy evaluation, and finite population size. Evolutionary game theory (EGT) thereby provides an elegant framework within which to study the behaviors of the various selection methods that can be used in a coevolutionary algorithm.

The canonical selection method used in evolutionary algorithms has individuals replicate (i.e., create offspring) in proportion to their evaluation score. This selection scheme corresponds to a standard replicator dynamic studied in EGT [3] and population dynamics [4]. The game-theoretic properties of the proportional-replication dynamic are well known [4]–[7]; in particular, all point attractors of proportional selection are Nash equilibria. Yet, many alternative selection methods are also used in coevolutionary algorithms, and these methods have escaped careful game-theoretic scrutiny. This paper aims to provide such an examination.

Manuscript received August 11, 2004; revised January 18, 2005 and June 2, 2005. The work of S. G. Ficici was done while with the Department of Computer Science, Brandeis University, Waltham, MA. This work was supported in part by the U.S. Department of Energy (DOE) under Grant DE-FG02-01ER45901.

S. G. Ficici is with the Artificial Intelligence Research Group, Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138 USA (e-mail: sevan@eecs.harvard.edu).

O. Melnik is with the Center for Discrete Mathematics and Theoretical Computer Science (DIMACS), Rutgers University, Piscataway, NJ 08854 USA (e-mail: melnik@dimacs.rutgers.edu).

J. Pollack is with the Department of Computer Science, Brandeis University, Waltham, MA 02454 USA (e-mail: pollack@cs.brandeis.edu).

Digital Object Identifier 10.1109/TEVC.2005.856203

We investigate the properties of several common selection methods from a game-theoretic and dynamical-systems perspective; the selection methods we discuss are proportional selection, truncation, $(\mu, \lambda), (\mu + \lambda)$, linear ranking, Boltzmann, and two types of tournament selection. These alternative selection methods have been reviewed and studied extensively in the context of "ordinary" evolutionary algorithms (i.e., non*co*evolutionary) [8]–[12]. But, how do these methods behave (and compare with standard proportional selection) in the context of a simple game-theoretic framework (and coevolution)? Do they exhibit similar dynamics and promote the same fixed-points and attractors?

Our investigation of selection dynamics employs simple variable-sum games whose Nash equilibria require mixed strategies, expressed as polymorphic populations. We focus on the behavior of selection methods around polymorphic equilibria for three reasons (all made apparent in this paper). First, polymorphic equilibria dramatically contrast the operation of the different selection methods. Second, polymorphic equilibria elicit nonintuitive and complex dynamics from most of the selection methods. Third, a game is easily constructed to have polymorphic equilibria. While proportional replication causes our simplified coevolutionary system to converge onto Nash equilibrium in our games, we find a variety of other behaviors from the alternative selection methods we list above. Rather than maintain the dynamics and Nash-equilibrium attractors of proportional selection, the alternative selection methods we test yield cyclic dynamics, non-Nash attractors, or chaos. Only Boltzmann selection (and only at low selection pressures) faithfully reproduces the attractors of proportional selection.

The conclusions to draw from our results hinge upon the purpose for which the coevolutionary algorithm is constructed. There are at least three ways in which to consider the operation of a coevolutionary algorithm. We can view the coevolutionary algorithm as: 1) a search method (i.e., a problem solver); 2) a model of some other dynamical system; or 3) a dynamical system in its own right. The results we present in this paper have relevance and utility for each of these viewpoints. With respect to search, our results show how many commonly used selection methods can prevent conventional coevolutionary algorithms from solving problems, where the solution is a polymorphic Nash equilibrium. With respect to modeling, our results emphasize the degree to which the model's behavior is sensitive to implementation details regarding selection—details that we might not otherwise believe to be critical. With respect to dynamical systems, our results provide a novel view into the behavior of selection methods that reveals ties to well-known one-dimensional (1-D) maps.

We begin with an overview of EGT in Sections II–VI, with the latter detailing the dynamics of the standard proportional replicator. Sections VII–XIII examine various alternatives to proportional selection. Section XIV considers our results with respect to two prominent uses of a coevolutionary algorithm: problem solving and modeling. In Section XV, we turn to the literature and present case studies, where alternative selection methods are used in a coevolutionary setting. Section XVI discusses future extensions to our analyses and Section XVII offers a summary of our findings and concluding remarks. We include Appendices A–C to review the concept of the preimage, the calculation of a dynamical system's Liapunov exponent, as well as a standard technique for analyzing the stability of multidimensional differentiable maps. This paper expands our earlier work [13], primarily by examining additional selection methods, defining all selection methods in equational form, and providing more extensive analysis and discussion.

II. Evolutionary Game Theory (EGT)

In conventional game theory, the strength of the Nash equilibrium concept rests upon three vital assumptions: 1) all agents are rational; 2) all agents have complete knowledge of their domain (i.e., the strategies available to all players and the payoffs); and 3) the agents have common knowledge of these assumptions (including this one). The stringency of these assumptions leaves Nash equilibrium vulnerable to criticism because we can easily imagine situations where any or all of these three assumptions cannot realistically hold. Particularly, the assumptions of conventional game theory are not met by most (if not all) biological agents.

The central achievement of *evolutionary* game theory [3], [14] is the introduction of an alternative motivation for the Nash equilibrium concept—one that is liberated from the strong requirements of rationality, knowledge of the game, and common knowledge. In EGT, these requirements are replaced by the principle of Darwinian evolution. A population of agents can achieve Nash equilibrium through a process of natural selection.

While EGT does away with certain assumptions, it makes several new ones. The standard EGT framework is as follows.

1) We have an infinitely large population of agents, each of which plays some pure strategy of the game being investigated.
2) There is *complete mixing*; that is, every agent interacts with every other agent and accumulates payoffs as it goes.
3) Game outcomes (payoffs) are noiseless.
4) After agent interaction is complete, agents reproduce in proportion to their cumulative payoffs; reproduction is asexual and without variation, hence, offspring are clones of their parents. This latter assumption implies that strategies absent from the initial population cannot appear later in time.

The assumptions made by EGT are, of course, idealizations. Nevertheless, one may argue that some natural systems sufficiently approximate these assumptions to make the EGT formalism useful; for other systems, the gap between reality and EGT's idealizations is too wide, and the strong assumptions of EGT need to be moderated in some way. The effects of weakened assumptions on the behavior of the EGT framework have been investigated (e.g., [15] and [16]); we discuss some of this work in detail in Section XV.

The next sections present the EGT framework in detail.

III. 2×2 Symmetric Games

We are interested in the behavior of selection methods with respect to polymorphic Nash equilibria. The simplest games that can have polymorphic equilibria are 2×2 symmetric vari-

able-sum games for two players. The simplicity of these games permits a clear exposition of selection dynamics without a loss of generality for the results we present. (We discuss games with more than two strategies and asymmetric games in Section XVI and in follow-up work, in preparation).

A generic two-strategy payoff matrix for a symmetric variable-sum game is shown in (1). By convention, the matrix entries are payoffs to the row player; thus, to obtain a player's payoff, we treat that player as the row player and the other as the column player. For example, if Player 1 uses strategy X and Player 2 uses strategy Y, then the payoff received by Player 1 is b; the payoff for Player 2 is c

$$\mathbf{G} = \begin{array}{c|cc} & \mathrm{X} & \mathrm{Y} \\ \hline \mathrm{X} & a & b \\ \mathrm{Y} & c & d \end{array} \tag{1}$$

$$w_{\mathrm{X}} = ap + b(1-p)$$
$$w_{\mathrm{Y}} = cp + d(1-p). \tag{2}$$

As we state above, we assume an infinitely large population of agents. Each agent in our infinite population plays one or the other pure strategy—mixed strategies (stochastic behaviors based upon probability distributions over the pure strategies) are not allowed. Instead, we express a mixed strategy by the distribution of pure strategies in the population; when a population contains more than one of the game's pure strategies, then the population is *polymorphic*. Since we have only two pure strategies, we can represent the state of our infinite population with a single variable p, which we will use to denote the proportion of X-strategists in our population. The proportion of Y-strategists is, therefore, $1-p$.

Given a population state p, the linear equations in (2) calculate the cumulative payoffs w_{X} and w_{Y}, received by X- and Y-strategists, respectively (after complete mixing). We can easily see that, as $p \to 0$, the cumulative score $w_{\mathrm{X}} \to b$ and $w_{\mathrm{Y}} \to d$; as $p \to 1$, the cumulative score $w_{\mathrm{X}} \to a$ and $w_{\mathrm{Y}} \to c$.

IV. Polymorphic Equilibrium

There exist exactly two payoff structures for the 2×2 matrix (1) that create a polymorphic Nash equilibrium. Case 1 is where $a > c$ and $b < d$ and Case 2 is where $a < c$ and $b > d$; note that we move from one case to the other by simply swapping rows of the payoff matrix (without swapping the columns). For reasons that are made clear below, we are interested only in Case 2.

Clearly, Case 2 contains an uncountable infinity of games. Further, if we select payoff values from a continuous random variable, such that the four payoffs are independent and identically distributed (i.i.d.), then Case 2 will occur exactly with probability 0.25 ($\Pr(a < c) \times \Pr(b > d) = 0.5^2 = 0.25$). Thus, the games of interest in this paper are not rare (polymorphic equilibria become even easier to obtain in larger games).

The payoff structure of Case 2 implies that the two lines described by (2) must intersect at some value of p, where $0 < p < 1$. We know this because $w_{\mathrm{X}} = b > d = w_{\mathrm{Y}}$ at $p = 0$ and $w_{\mathrm{X}} = a < c = w_{\mathrm{Y}}$ at $p = 1$; therefore, $w_{\mathrm{X}} = w_{\mathrm{Y}}$ somewhere in-between. We call this intersection a *polymorphic equilibrium*; "polymorphic" because the population contains more than one of the game's pure strategies (which, for our 2×2 games means that $0 < p < 1$), and "equilibrium" because all of the agents in the population (and, hence, all the strategies that are being played in the population) receive the same cumulative score. Fig. 1 illustrates, without loss of generality, an example of polymorphic equilibrium.

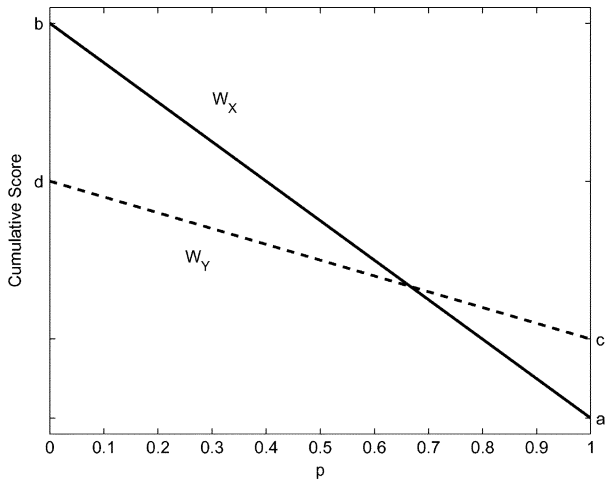

Fig. 1. Payoff structure of Case 2, where $a < c$ and $b > d$, causes w_{X} to intersect w_{Y} for some value of p, where $0 < p < 1$. (Figure adapted from [3, p. 181].)

A. Properties

We denote the proportion at which polymorphic equilibrium is achieved as p_{EQ}. To calculate the value $p = p_{\mathrm{EQ}}$, we set $w_{\mathrm{X}} = w_{\mathrm{Y}}$ and solve for p; the general form of the result is given by (3). Case 2 has the following three properties: 1) it can have no more than a single polymorphic equilibrium (because the lines can intersect no more than once); 2) p_{EQ} lies in the open interval (0, 1); and 3) for each possible value of p_{EQ} there exists an uncountable infinity of games that map to it—a) the value of p_{EQ} is unchanged if we add a constant to all four payoffs (we denote this constant as w_0, below) and b) we can rotate the lines w_{X} and w_{Y} about their intersection at p_{EQ}, keeping the slope of w_{X} less than that of w_{Y}, and obtain different games that share the same value of p_{EQ}. Properties 3a) and 3b) are easy to verify both algebraically and visually

$$p_{\mathrm{EQ}} = \frac{d-b}{a-c+d-b}. \tag{3}$$

Away from polymorphic equilibrium (i.e., for $p \neq p_{\mathrm{EQ}}$), the cumulative payoff for one strategy must be greater than that for the other. Below the equilibrium proportion (i.e., for all $p < p_{\mathrm{EQ}}$), we see that $w_{\mathrm{X}} > w_{\mathrm{Y}}$. Therefore, when placed in an evolutionary framework, all population states $p < p_{\mathrm{EQ}}$ will be acted upon by selection such that strategy X is favored, causing the value of p to increase. A similar statement can be made for population states $p > p_{\mathrm{EQ}}$, except that now $w_{\mathrm{X}} < w_{\mathrm{Y}}$; selective pressure favors strategy Y, causing p to decrease. Thus, we may intuit that the effect of selection will be similar to that of a negative feedback loop—p always moves toward p_{EQ}. Negative feedback is associated with stable dynamics and homeostasis and is what gives Case 2 special interest. Case 1 has all

the properties of Case 2, except that it creates a positive feedback loop—p always moves toward zero or one, depending upon which side of p_{EQ} the initial population state lies. (Case 1 is illustrated by Fig. 1 by swapping the line labels.)

In Case 1 (positive feedback), each selection method that we examine behaves as we intuit—the population converges to $p = 0$ or $p = 1$, depending upon the initial condition. In Case 2 (negative feedback), however, we find that only proportional selection (and Boltzmann selection at sufficiently low selection pressures) converges onto p_{EQ} as we intuit; most alternatives to proportional selection are actually made unstable by the negative feedback loop and fail to converge onto p_{EQ}. The sections below examine precisely why this is the case.

B. Nash Equilibrium

We know that p_{EQ} represents a state of score-equilibrium between X- and Y-strategists. The proportion of X- and Y-strategists at p_{EQ} turns out also to represent a *Nash equilibrium* [17], [18] of the game (this is true for both Cases 1 and 2). A Nash-equilibrium strategy for a symmetric game is a (possibly mixed) strategy s^* that is its own *best reply*. That is, if one player uses s^*, then the highest payoff that the other player can obtain is received by also using s^*; if both players use s^*, then neither has incentive to deviate unilaterally from that strategy and so the players are in Nash equilibrium. Though our agents are strictly pure strategists, we can interpret a polymorphic population, as a whole, to represent a mixed strategy; particularly, our polymorphic equilibrium indicates a mixed Nash-equilibrium strategy, where the probability of playing strategy X is p_{EQ} (and $\Pr(\mathrm{Y}) = 1 - p_{\mathrm{EQ}}$).

For the purpose of clarity, we focus our attention in this paper on polymorphic Nash equilibria that involve only two strategies. Larger symmetric variable-sum games may have polymorphic Nash equilibria that involve more than two strategies. Regardless of the number of strategies involved, all strategies played in a polymorphic Nash equilibrium will be at fitness equilibrium [3], [19].

C. Domination

Above we consider the two payoff structures that yield a polymorphic equilibrium. When a 2×2 symmetric game lacks a polymorphic equilibrium, then one of the game's two strategies *dominates* the other. Domination is achieved either when $a > c$ and $b > d$, where strategy X dominates strategy Y, or when $a < c$ and $b < d$, where Y dominates X (for either case, we may substitute an inequality with equality, giving a weaker form of domination).

Without loss of generality, let us assume $a > c$ and $b > d$. When strategy X dominates Y, we obtain $w_{\mathrm{X}} > w_{\mathrm{Y}}$ regardless of the population state p; thus, w_{X} and w_{Y} cannot intersect and we cannot have a polymorphism. This payoff structure makes strategy X the Nash-equilibrium strategy of the game; it also creates a positive feedback loop, making $p = 1.0$ the attractor (in our simple framework) for all of the selection methods we discuss in this paper. Fitness equilibrium (amongst agents) is achieved only when all players are X-strategists.

Selection dynamics in games with domination are similar to the selection dynamics obtained outside of coevolution—that is, in "ordinary" evolutionary algorithms that are used to optimize static objective functions. The rating that an individual receives from the objective function does not depend upon what other individuals exist in the population. If genotype A is rated higher than genotype B, then A will eventually take over a population of A- and B-types (in the absence of variation), regardless of its initial frequency (>0); this is much like how a dominating strategy takes over in our framework.

Thus, whenever a positive feedback loop exists (with or without a polymorphic equilibrium) the eventual outcomes of all selection methods are identical in our framework. This common behavior obtained with positive feedback tempts us to expect similarly unified behavior in the case of polymorphisms with negative feedback. This paper shows that such an expectation is mistaken.

V. Hawk–Dove Game

To illustrate the dynamics of selection concretely, we must instantiate the payoff structure we discuss above (i.e., Case 2: $a < c$ and $b > d$). Fortunately, the Hawk–Dove game ([3, Ch. 2]) is a well-known and frequently studied game (e.g., [20, Ch. 5]) that can have a polymorphism with the negative-feedback structure described in Section IV. The generalized version of the game has two parameters, $V > 0$ (the *value* of a resource) and $C > 0$ (the *cost* incurred by fighting to gain control of the resource), shown in the payoff matrix (4). The Hawk–Dove game has two pure strategies, *Hawk* and *Dove*. Hawks are always willing to fight, and Doves are never willing to fight. Thus, two Hawks incur the cost of fighting, whereas two Doves divide the resource peacefully with each other; Doves yield entirely to Hawks. (See [3, Ch. 2] for Maynard Smith's detailed rationalization for this parameterization.)

If $V > C$, then the game's payoff structure is $a > c$ and $b > d$. The Hawk strategy dominates Dove, and $w_{\mathrm{H}} > w_{\mathrm{D}}$ regardless of the population state p. If $V < C$, then the payoff structure of the Hawk–Dove game creates, without loss of generality, the polymorphism with negative feedback that we discuss in Section IV. The equilibrium (3) reduces to $p_{\mathrm{EQ}} = V/C$

$$\mathbf{G} = \begin{array}{c|cc} & \text{Hawk} & \text{Dove} \\ \hline \text{Hawk} & (V-C)/2 & V \\ \text{Dove} & 0 & V/2 \end{array}. \tag{4}$$

Since the standard EGT framework has agents reproduce in proportion to cumulative payoffs, and an agent cannot have fewer than zero offspring, we must ensure that cumulative payoffs are nonnegative. For this reason, we make all payoffs in the matrix greater than zero by adding to each of them a sufficiently large constant w_0. Recall that such a constant does not alter the value of p_{EQ} (see Section IV-A).

Before we proceed with our analyses, we wish to emphasize the generality of our results. A quick inspection of the replicator equations we examine in this paper reveals that only proportional (6) and Boltzmann (20) selection make explicit reference to the payoffs a, b, c, and d. In contrast, none of the other methods we survey uses the payoff values in itsr calculations—only the current state p and the equilibrium state p_{EQ} are used to calculate the new population state $f(p)$. As

we note in Section IV-A, each value of p_{EQ} is mapped to by an uncountable infinity of games; by not referring to the actual payoff values, these other selection methods cannot distinguish between two different games that share the same equilibrium p_{EQ}. Thus, for each of the selection methods other than proportional and Boltzmann, two different games that share the same equilibrium point p_{EQ} will yield *precisely identical behavior* in all respects. For proportional selection, two different games that share the same equilibrium point p_{EQ} will yield identical qualitative behavior, but differ in detail (they may have different rates of convergence); Botlzmann selection can also yield the same qualitative behavior between the games, but may require different selection pressures be used. For these reasons, we can use the Hawk–Dove game (with $V < C$) to derive the results in this paper without loss of generality. Section XVI discusses how our results generalize further to higher-dimensional spaces. (To simplify our discussion, we specify games with V, C, and w_0 for proportional and Boltzmann selection only; otherwise, we state only the value of p_{EQ} to specify a game.)

VI. Score-Proportional Selection

As we state above, agents in the conventional EGT framework reproduce in proportion to their cumulative payoffs. For the general two-strategy payoff-matrix in (1), the cumulative payoff equations are given by (2). Thus, given a population state p (representing the proportion of X-strategists in the population), the proportion of X-strategists in the next generation $f(p)$ is given by (5). This difference equation is known as a *discrete-time replicator* in EGT parlance; dynamical-systems theory refers to such an equation as a *map*. After simplification, we obtain (6). The derivative of the replicator is given in (7)

$$f(p) = \frac{pw_{\mathrm{X}}}{pw_{\mathrm{X}} + (1-p)w_{\mathrm{Y}}} \tag{5}$$

$$f(p) = \frac{p^2(a-b)+pb}{p^2(a-b-c+d)+p(b+c-2d)+d} \tag{6}$$

$$f'(p) = \frac{p^2(bd-2ad+ac)+p(2ad-2bd)+bd}{(p^2(a-b-c+d)+p(b+c-2d)+d)^2}. \tag{7}$$

We can visualize the dynamics of (6) with a *map diagram*. Fig. 2 shows the map diagram generated by proportional selection on the Hawk–Dove game, where $p_{\mathrm{EQ}} = 2/3 (V = 8, C = 12, w_0 = 2.005)$. The x axis represents our population state p at time t; the y axis represents the population state at time $t+1$, that is $f(p)$. The bold curve represents the behavior of our map. Where this curve intersects the diagonal, we have a *fixed-point* in our map, which is a population state p such that $f(p) = p$; that is, a fixed-point is a population state that does not change from one time-step to the next. Our map has three fixed-points. Two of them are $p = 0$ and $p = 1$; these are *monomorphic* population states that contain only Doves and only Hawks, respectively. The third fixed-point is polymorphic and occurs at p_{EQ}. As we know from Fig. 1, all agents (Hawks and Doves) at p_{EQ} obtain the same cumulative score; therefore, score-proportionate reproduction leaves the population state unchanged. But, what happens over time when $p \neq p_{\mathrm{EQ}}$?

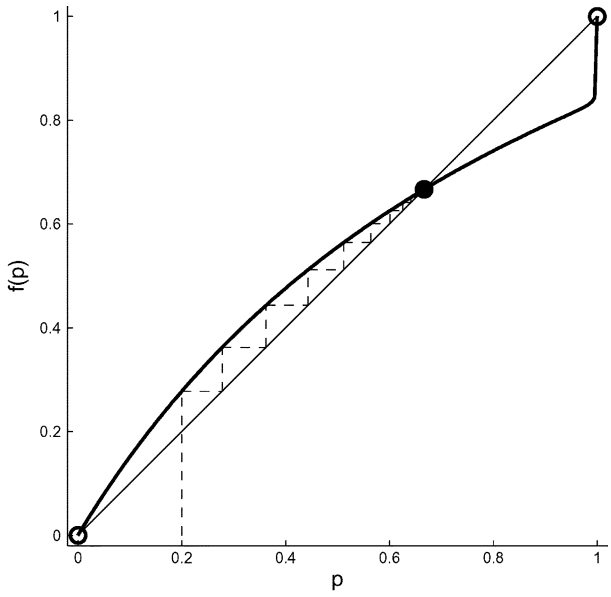

Fig. 2. Map diagram of score-proportional selection operating on Hawk–Dove game, where $p_{\mathrm{EQ}} = 2/3$ $(v = 8, c = 12, w_0 = 2.005)$. Orbit has initial condition $p_{\mathrm{IC}} = 0.2$ and converges onto $p_{\mathrm{EQ}} = 2/3$. Stable fixed-point is denoted by filled circle; unstable fixed-points are denoted by open circles. To read an orbit on a map diagram, begin on the x axis at $(p_{\mathrm{IC}}, 0)$, proceed up to the curve at $(p_{\mathrm{IC}}, f(p_{\mathrm{IC}}))$, then alternately move horizontally to the diagonal and vertically to the curve.

If the derivative of a 1-D map at a fixed-point has an absolute value less than one, then the fixed-point is *stable* [21]; if the derivative's absolute value is greater than one, then the fixed-point is unstable (see Appendix C for details). A quick visual inspection of Fig. 2 reveals that p_{EQ} is stable, whereas $p = 0$ and $p = 1$ are unstable. This means that if we perturb a population at p_{EQ}, then the population will converge back onto p_{EQ}. In contrast, if we perturb a population composed entirely of Doves ($p = 0$), by adding any amount of Hawks, then the population state will diverge from $p = 0$ (and eventually converge onto p_{EQ}); a similar statement applies to a population of all Hawks ($p = 1$). Thus, p_{EQ} is the unique, global *attractor* of our proportional-reproduction dynamical system; any initial population state $0 < p_{\mathrm{IC}} < 1$ asymptotically converges onto p_{EQ}. The dashed line in Fig. 2 shows an example trajectory, or *orbit*, that our population can take over time, here starting from an initial condition of $p_{\mathrm{IC}} = 0.2$.

More generally, proportional selection is known to have the following properties [4]. All Nash equilibria are fixed-points of proportional selection (6). All point attractors of proportional selection are Nash equilibria. Not all Nash equilibria are attractors, however (for example, the Nash equilibrium of the familiar Rock-Paper-Scissors game is not a point-attractor). Indeed, using a continuous-time version of proportional selection, Skyrms [22] illustrates a four-strategy symmetric game that creates chaotic dynamics (and as Skyrms points out, continuous-time replication dynamics are typically better behaved than their discrete-time counterparts).

Having established the equilibria and dynamics of the standard replicator for two-strategy games that have polymorphic attractors, we now begin our examination of the replication dynamics generated by several alternative selection methods.

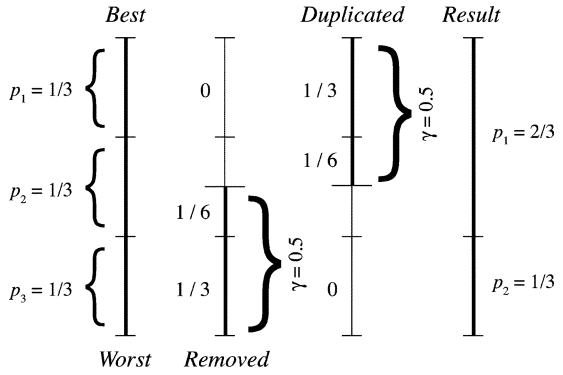

Fig. 3. Operation of truncation selection in an infinite population. (Figure adapted from [13].)

VII. Truncation Selection

Truncation selection is so called because individuals below some quality threshold are removed (i.e., "truncated") from the pool individuals that will be used to create offspring. The threshold can express an absolute quality level or a relative one. Truncation is often used in the branch of evolutionary computation known as *evolutionary programming* [23].

The particular form of truncation selection we study here is found, for example, in [16], [24], and proceeds as follows: 1) we sort the population according to the agents' evaluation scores; 2) we remove the worst γ fraction of the population; and 3) we replace the removed agents with variations of individuals in the best γ fraction of the population. For example, let us say we have a population of 200 agents and $\gamma = 0.25$. In this case, each of the best 50 agents will create one offspring, which will in turn replace one of the worst 50 agents; the top 75% of the population is left unaltered.

Since our framework lacks variation operators, offspring are clones of their parents. In this case, selection pressure is maximized at $\gamma = 0.5$ and is reduced as $\gamma \to 0$; values of $\gamma = 0.5+x$ (for $0 \leq x \leq 0.5$) behave exactly as $\gamma = 0.5 - x$. When we extend γ beyond 0.5, then the set of agents that create offspring must overlap with the set of agents that are replaced. Without loss of generality, we can assume that each agent in that overlapping subset is replaced by its own identical offspring (as opposed to another agent's offspring). Thus, the overlapping set of agents has no net effect on the change in population-state.

Truncation selection is easily implemented in our infinite-population EGT framework. Since all agents that play the same strategy receive the same fitness, we only need to sort the strategies and note the proportion with which each strategy appears in the population. Fig. 3 gives an example where our population has three pure strategies $s_{1,\ldots,3}$ played with proportions $p_{1,\ldots,3}$, respectively; we set $\gamma = 0.5$, and we begin with a population where $p_i = 1/3$ for $i = 1, \ldots, 3$. The strategies are sorted by their cumulative scores; the highest score belongs to s_1 and lowest to s_3 (left-most column). Since $\gamma = 0.5$, we remove the worst 1/2 of the population (second column); this includes all of s_3 and 1/2 of s_2. We then duplicate all of s_1 and 1/2 of s_2 (third column) to arrive at our new population (fourth column), where s_1 and s_2 now have $2/3$ and $1/3$ of the population, respectively; s_3 is eliminated.

For 1-D systems (i.e., symmetric games with just two strategies), the precise operation of truncation selection is specified by the map (8); the derivative of the truncation map is given by (9). If $p < p_{\mathrm{EQ}}$, then Hawks outscore Doves, and we increase the proportion of Hawks. The first three cases in (8) concern the magnitude of this increase (by symmetry, the latter three cases concern the situation where $p > p_{\mathrm{EQ}}$ and Doves outscore Hawks). Clearly, a single application of the truncation map can no more than double the proportion of Hawks (Case 3). Further, if the proportion of Hawks is greater than γ, then we merely increase p by γ (Case 2), rather than double p. Finally, we must limit p to no more than 1.0 (Case 1). Note that the operation of truncation selection depends upon the game's payoffs only to the extent that they determine the value of p_{EQ}; once the equilibrium value is determined, the payoffs are irrelevant to the selection dynamics

$$f(p) = \begin{cases} 1.0, & \text{if } p < p_{\mathrm{EQ}} \text{ and } p \geq 1-\gamma \\ p+\gamma, & \text{if } p < p_{\mathrm{EQ}} \text{ and } \gamma \leq p < 1-\gamma \\ 2p, & \text{if } p < p_{\mathrm{EQ}} \text{ and } p < \gamma \\ p-\gamma, & \text{if } p > p_{\mathrm{EQ}} \text{ and } \gamma < p \leq 1-\gamma \\ 2p-1, & \text{if } p > p_{\mathrm{EQ}} \text{ and } p > 1-\gamma \\ 0, & \text{if } p > p_{\mathrm{EQ}} \text{ and } p \leq \gamma \\ p_{\mathrm{EQ}}, & \text{if } p = p_{\mathrm{EQ}} \end{cases} \tag{8}$$

$$f'(p) = \begin{cases} 0, & \text{if } p < p_{\mathrm{EQ}} \text{ and } p \geq 1-\gamma \\ 1, & \text{if } p < p_{\mathrm{EQ}} \text{ and } \gamma \leq p < 1-\gamma \\ 2, & \text{if } p < p_{\mathrm{EQ}} \text{ and } p < \gamma \\ 1, & \text{if } p > p_{\mathrm{EQ}} \text{ and } \gamma < p \leq 1-\gamma \\ 2, & \text{if } p > p_{\mathrm{EQ}} \text{ and } p > 1-\gamma \\ 0, & \text{if } p > p_{\mathrm{EQ}} \text{ and } p \leq \gamma \\ \text{undefined}, & \text{if } p = p_{\mathrm{EQ}} \end{cases} . \tag{9}$$

The truncation map is piecewise linear with a discontinuity at p_{EQ}. If $p_{\mathrm{EQ}} < \gamma$, then we can converge onto the monomorphic fixed-point $p = 0$; a variety of unstable cycles are also possible. If $p_{\mathrm{EQ}} > 1-\gamma$, then we can converge onto the monomorphic fixed-point $p = 1$; a variety of unstable cycles are also possible. If $\gamma < p_{\mathrm{EQ}} < 1-\gamma$, then we get either neutrally-stable two-cycles or chaos. The chaotic regime requires $1 - 2\gamma < p_{\mathrm{EQ}} < 2\gamma$; this regime also contains unstable cycles. We review these regimes of behavior below.

A. Regime I: Point Attractors and Unstable Cycles

The truncation map yields a great variety of possible behaviors. We can easily imagine one example with a Hawk–Dove game where $p_{\mathrm{EQ}} = 7/12$. Assume $\gamma = 0.5$ and a population state $0.5 \leq p < 7/12$. Since $p < p_{\mathrm{EQ}}$, we know that Hawks outscore Doves; since $p \geq 0.5$, we also know that Hawks are at least half of the population. Thus, the best half is only Hawks. Replacing the worst half, we arrive at a population of all Hawks. What happened to our Hawk–Dove polymorphism?

Fig. 4 (top) shows the map diagram for the above example. The most noticeable feature is the large discontinuity at p_{EQ}—precisely where our polymorphic Nash-equilibrium attractor should be; p_{EQ} is now a singularity and unstable. For most initial conditions, such as the one shown, the orbit will eventually enter the critical interval $[1/2, 7/12)$, whereupon convergence to $p = 1$ (a monomorphic fixed-point) is guaranteed. Note that convergence to $p = 1$ is not asymptotic;

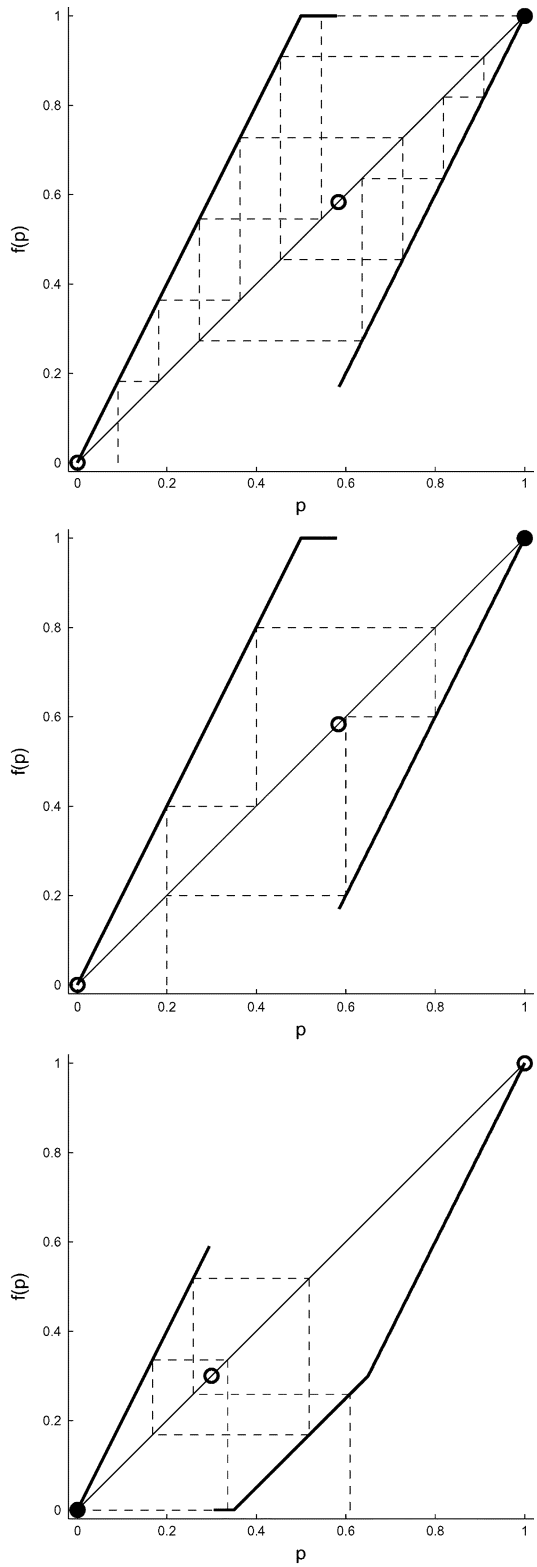

Fig. 4. Truncation map with Hawk–Dove game. Top: Orbit converges to $p = 1.0$ ($p_{\rm EQ} = 7/12, \gamma = 0.5, p_{\rm IC} = 1/11$). Middle: Orbit enters unstable period-four cycle ($p_{\rm EQ} = 7/12, \gamma = 0.5, p_{\rm IC} = 1/5$). Bottom: Orbit converges to $p = 0$ ($p_{\rm EQ} = 0.3, \gamma = 0.35, p_{\rm IC} = 0.609$).

in particular, if we slightly perturb a population at $p = 1$, then it will move away from this state before returning via the $[1/2, 7/12)$ interval. Thus, $p = 1$ is an attractor, but it is not locally stable. In general, convergence to $p = 1$ is possible

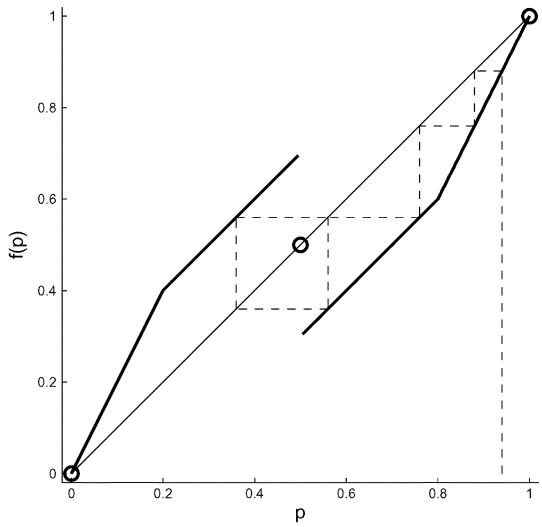

Fig. 5. Orbit enters neutrally stable cycle ($p_{\rm EQ} = 0.5, \gamma = 0.2, p_{\rm IC} = 0.94$).

when $p_{\rm EQ} > 1 - \gamma$; the critical interval for which $f(p) = 1$ is $1 - \gamma \leq p < p_{\rm EQ}$.

Not all initial conditions converge to $p = 1$, however; exceptions produce unstable cycles or the unstable fixed-point $p_{\rm EQ}$. For example, the initial conditions $p_{\rm IC} = 1/3, 5/7$, and $1/5$ give unstable two-, three-, and four-cycles, respectively; Fig. 4 (middle) shows the latter case. If $p_{\rm IC}$ is in the iterated preimage of $p_{\rm EQ}$ (see Appendix A), then the orbit enters the fixed-point $p_{\rm EQ}$. Let $\mathcal{S}$ be the set of initial conditions that enter a particular unstable outcome (a cycle or $p_{\rm EQ}$). The elements of $\mathcal{S}$ are clearly both *disconnected* (i.e., do not form a continuum) and *isolated* (i.e., do not have a neighbor arbitrarily nearby). Therefore, virtually any perturbation will push an orbit (with $p_{\rm IC} \in \mathcal{S}$) out of $\mathcal{S}$, cause it to diverge from the cycle (or $p_{\rm EQ}$), and eventually converge onto $p = 1.0$.

By symmetry, if $p_{\rm EQ} < \gamma$, then the truncation map has an attractor at $p = 0$, and most initial conditions converge onto it. This possibility is illustrated by Fig. 4 (bottom). The critical interval in this case is $(p_{\rm EQ}, \gamma]$. Of course, unstable cycles are also possible, as above. Because selection pressure must be in the range $0 \leq \gamma \leq 0.5$, it must be the case that $\gamma \leq 1 - \gamma$; thus, the attractors $p = 0$ and $p = 1$ are mutually exclusive (we cannot have $1 - \gamma < p_{\rm EQ} < \gamma$).

B. Regime II: Neutrally Stable Cycles

If $\gamma < p_{\rm EQ} < 1 - \gamma$, then yet newer possibilities arise. Fig. 5 shows the truncation map when $p_{\rm EQ} = 0.5$ and $\gamma = 0.2$. Here, all initial conditions $0 < p_{\rm IC} < 1$ (except for the preimages of $p_{\rm EQ}$, of course), enter neutrally stable two-cycles, as the example orbit demonstrates. The cycle begins with the first pair of consecutive population states $\langle p, f(p) \rangle$ that are on opposite sides of $p_{\rm EQ}$. To this pair, we can add or subtract any ϵ (that keeps p and $f(p)$ on opposite sides of $p_{\rm EQ}$) and the new pair will remain in the shifted two-cycle. The amplitude of the cycle $|f(p) - p|$ is equal to γ. Thus, as selection pressure is decreased ($\gamma \to 0$), cyclic trajectories will orbit closer and closer to the score equilibrium $p_{\rm EQ}$. Note that $p_{\rm EQ}$ remains unstable.

C. Regime III: Chaos

If $1-2\gamma < p_{\mathrm{EQ}} < 2\gamma$ (which is a tighter constraint on γ than $\gamma < p_{\mathrm{EQ}} < 1-\gamma$), then truncation selection permits chaos. So long as this constraint holds, the ease with which chaos can be obtained increases with γ. A truncation map that produces chaos may still have a fairly large domain that yields neutrally stable cycles; Fig. 6 (top) illustrates such a cycle (the Liapunov exponent λ is zero). Fig. 6 (middle) shows the same map producing a chaotic orbit. Fig. 6 (bottom) gives the histogram of population states visited in the chaotic orbit. The Liapunov exponent for this chaotic orbit is $\lambda \approx 0.346\,57$; this quantity is approximately $\ln(2)/2$, which reflects the fact that approximately half of the orbit touches the map segments with slope 2.0, while the rest of the orbit touches segments with slope 1.0.

The salient change that allows chaos concerns the portions of the map that have slope 2.0. In each of our previous examples, these regions are visited only as part of an initial transient or unstable cycle (if at all). In a chaotic orbit, these segments of the map are revisited throughout, and are responsible for giving $\lambda > 0$, which is an indicator of chaos. (See Appendix B for details on the Liapunov exponent).

The histogram shows the chaotic orbit to jump between three distinct regions in state space. The two gaps $p \in [0.3, 0.4]$ and $p \in [0.6, 0.7]$, which separate these three regions, are the portions of state space that support the neutrally stable cycle. Once an orbit enters one of these two gaps, it will exclusively visit portions of the map that have slope 1.0 (this is easily verified by examining the map equation). Any perturbation that keeps an orbit within these gaps allows the cyclic dynamic to continue.

Fig. 7 shows the truncation map with $p_{\mathrm{EQ}} = 0.5$ and $\gamma = 0.43$. The map segments with slope 2.0 are much longer, which results in a higher Liapunov exponent of $\lambda \approx 0.565\,19$. The orbit's distribution over the population states is approximately uniform over the interval [0.07, 0.93]. Given a game with $p_{\mathrm{EQ}} = 0.5$, as $\gamma \to 0.5$, the portions of the map with slope 1.0 vanish and $\lambda \to \ln(2) \approx 0.6931$. The map produced with $p_{\mathrm{EQ}} = 0.5$ and $\gamma = 0.5$ is identical to the well-known Bernoulli shift map (also known as the binary shift map) [25, Ch. 10.5].

VIII. (μ, λ) and $(\mu + \lambda)$ Selection

The (μ, λ) and $(\mu + \lambda)$ selection methods originate from the branch of evolutionary computation known as *evolution strategies* [26]–[28]. In both methods, we begin with a population of μ agents. To create the next generation, each of these agents produces k offspring (with variation), resulting in $\lambda = k\mu$ total offspring (the use of λ here is not to be confused with the Liapunov exponent). In (μ, λ) selection, the best μ agents from this larger set of λ offspring are selected to form the next generation of our evolving population; in $(\mu + \lambda)$ selection, the best μ agents from amongst the λ offspring and μ parents are chosen to form the next generation. We can think of $(\mu + \lambda)$ as (μ, λ) selection where each of the μ parents generates an additional offspring that is an exact copy.

Since we have an infinite population, the exact values of μ and λ are unimportant—we care only about the quotient μ/λ, which we will define to be our selection parameter $\gamma = \mu/\lambda$.

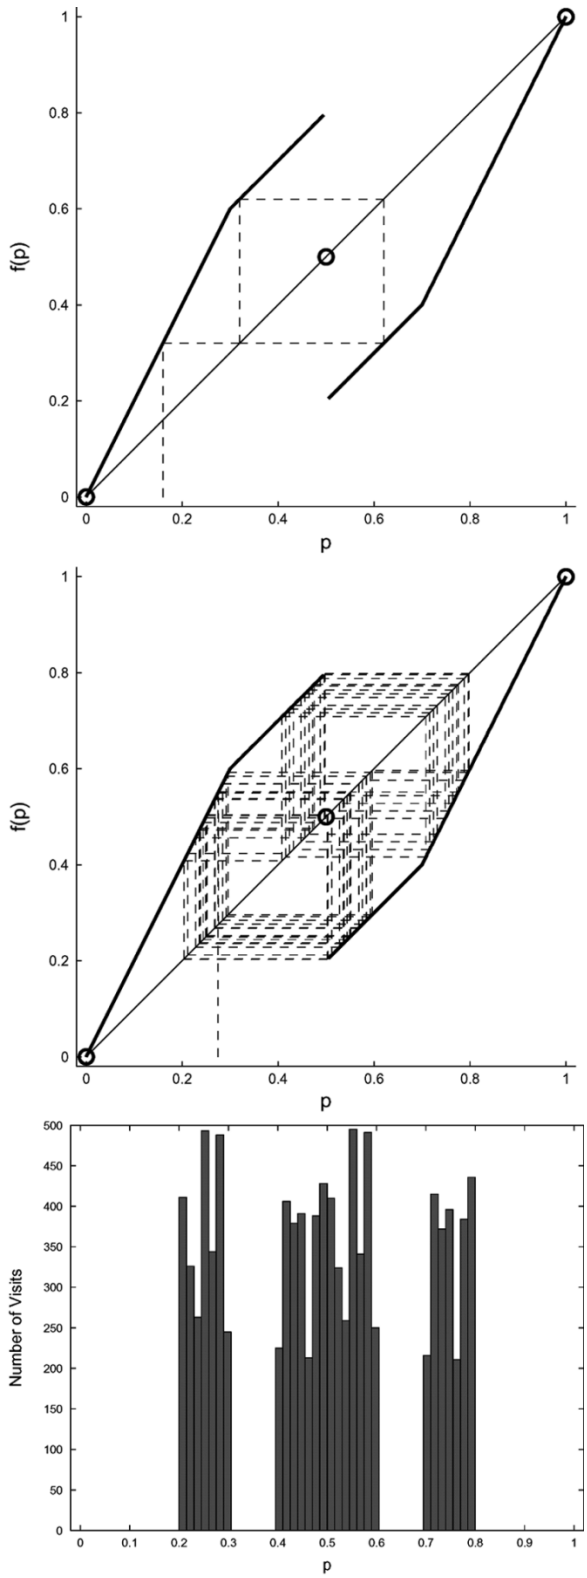

Fig. 6. Truncation map that permits neutral cycles and chaos ($p_{\mathrm{EQ}} = 0.5, \gamma = 0.3$). Top: Orbit that enters neutrally stable two-cycle ($p_{\mathrm{IC}} = 0.16$). Middle: First 100 time-steps of chaotic orbit ($p_{\mathrm{IC}} = 0.2745$). Bottom: Histogram of population states visited in chaotic orbit (10^4 time-steps, $p_{\mathrm{IC}} = 0.2745, \lambda \approx 0.346\,57$).

The value of γ lies in the interval $\gamma \in (0, 1)$ and represents the proportion of our original population that gets to have its offspring in the next generation. Smaller values of γ, therefore,

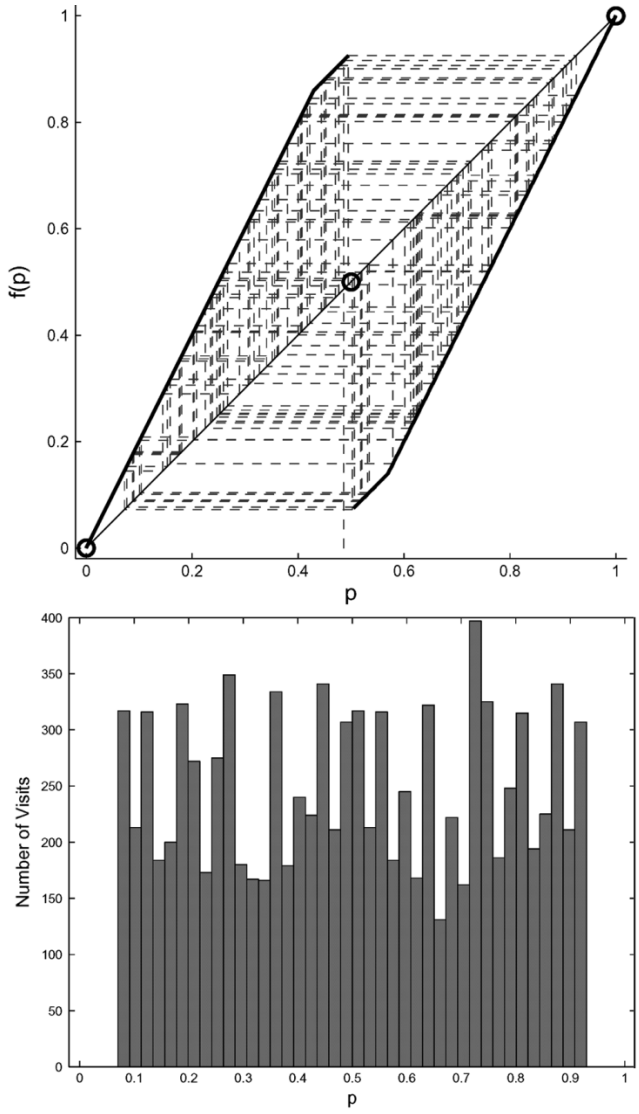

Fig. 7. Truncation map with stronger chaotic behavior ($p_{\mathrm{EQ}} = 0.5, \gamma = 0.43, p_{\mathrm{IC}} = 0.4859$). Top: First 100 time-steps of chaotic orbit. Bottom: Histogram of population states visited in chaotic orbit (10^4 time-steps, $\lambda \approx 0.565\,19$).

produce a stronger selection pressure. Because our framework generates offspring asexually and without variation, (μ, λ) and $(\mu + \lambda)$ selection operate identically; further, (μ, λ) selection (and thus, $(\mu + \lambda)$ selection) reduces to the more conventional (and drastic) form of truncation selection: agents above the selection threshold produce enough offspring to replace the entire population, rather than just the weaker individuals as in Section VII.

In our 1-D system, the equation that governs (μ, λ) is given by (10); the derivatives of this map are given by (11). The map equation is easily understood. In essence, we simply rescale the proportion of the higher-scoring strategy according to γ; the first and fourth cases of our map equation simply limit p to no more than one and no less than zero.

The (μ, λ) method produces a number of possible behaviors. Like truncation selection, the (μ, λ) map is piecewise linear, and we find a discontinuity at p_{EQ}; the Nash equilibrium is again a singularity and unstable. Note that, for any value of p_{EQ}, (μ, λ) with $\gamma = 0.5$ behaves identically to truncation (8) with $\gamma = 0.5$

$$f(p) = \begin{cases} 1, & \text{if } p < p_{\mathrm{EQ}} \text{ and } p \geq \gamma \\ p/\gamma, & \text{if } p < p_{\mathrm{EQ}} \text{ and } p < \gamma \\ 1 + (p-1)/\gamma, & \text{if } p > p_{\mathrm{EQ}} \text{ and } p > 1 - \gamma \\ 0, & \text{if } p > p_{\mathrm{EQ}} \text{ and } p \leq 1 - \gamma \\ p_{\mathrm{EQ}}, & \text{if } p = p_{\mathrm{EQ}} \end{cases} \quad (10)$$

$$f'(p) = \begin{cases} 0, & \text{if } p < p_{\mathrm{EQ}} \text{ and } p \geq \gamma \\ 1/\gamma, & \text{if } p < p_{\mathrm{EQ}} \text{ and } p < \gamma \\ 1/\gamma, & \text{if } p > p_{\mathrm{EQ}} \text{ and } p > 1 - \gamma \\ 0, & \text{if } p > p_{\mathrm{EQ}} \text{ and } p \leq 1 - \gamma \\ \text{undefined}, & \text{if } p = p_{\mathrm{EQ}} \end{cases} . \quad (11)$$

A. Regime I: Point Attractors and Unstable Cycles

If $p_{\mathrm{EQ}} > 1 - \gamma \geq \gamma$, then $p = 1$ becomes an attractor; if $p_{\mathrm{EQ}} < \gamma \leq 1 - \gamma$, then $p = 0$ becomes an attractor. Both of these attractors are simultaneously available if $\gamma < p_{\mathrm{EQ}} < 1 - \gamma$. Fig. 8 (top) shows a sample orbit that converges to $p = 0$.

Fig. 8 (bottom) shows that this map ($p_{\mathrm{EQ}} = 0.4, \gamma = 0.3$) also has an unstable two-cycle. We discover the location of the cycle by looking for fixed-points of the *second iterate* of our map; that is, $f^2(p_{\mathrm{Cycle}}) = p_{\mathrm{Cycle}}$, where $f^2(p) \equiv f(f(p))$. Since a cycle must cross p_{EQ}, we arbitrarily choose $p_{\mathrm{Cycle}} < p_{\mathrm{EQ}}$. Thus, a two-cycle must satisfy (12); solving for p_{Cycle}, we obtain (13). Given $\gamma = 0.3$, our equation gives us $p_{\mathrm{Cycle}} = 0.21/0.91 = 3/13 \approx 0.2308$; since $p_{\mathrm{Cycle}} < p_{\mathrm{EQ}}$ and $f(p_{\mathrm{Cycle}}) > p_{\mathrm{EQ}}$, all of our requirements are met. Because this cycle is unstable, any perturbation will cause it to diverge and eventually converge onto $p = 0$ or $p = 1$. (The instability of this cycle is easily demonstrated by the accumulation of round-off error in numerical simulation.)

$$f^2(p_{\mathrm{Cycle}}) = 1 + \frac{\frac{p_{\mathrm{Cycle}}}{\gamma} - 1}{\gamma} = p_{\mathrm{Cycle}} \quad (12)$$

$$p_{\mathrm{Cycle}} = \frac{\gamma^2 - \gamma}{\gamma^2 - 1}. \quad (13)$$

B. Regime II: Chaos

If $1 - \gamma < p_{\mathrm{EQ}} < \gamma$, then (μ, λ) produces either chaos or unstable cycles. Fig. 9 (top) illustrates a chaotic orbit where $p_{\mathrm{EQ}} = 0.5, \gamma = 0.6$, and $p_{\mathrm{IC}} = 0.65$. Since the slope of the map is $1/\gamma$ throughout, the Liapunov exponent is $\lambda = \ln(1/\gamma) \approx 0.510\,83$. Fig. 9 (bottom) shows the histogram of population states visited by the orbit. The multimodal distribution contrasts with the near-uniform distribution produced by truncation selection on the same game [Fig. 7 (bottom)]. The symmetry of Fig. 9 (bottom) follows from the symmetric location of score equilibrium. We can apply (13) to obtain an unstable two-cycle in this map, as well; the appropriate initial condition is $p_{\mathrm{IC}} = 0.375$. As γ increases, the slopes of the map approach 1.0 and converge onto the diagonal. As a result, the overall rate of state-change decreases, as does the Liapunov exponent.

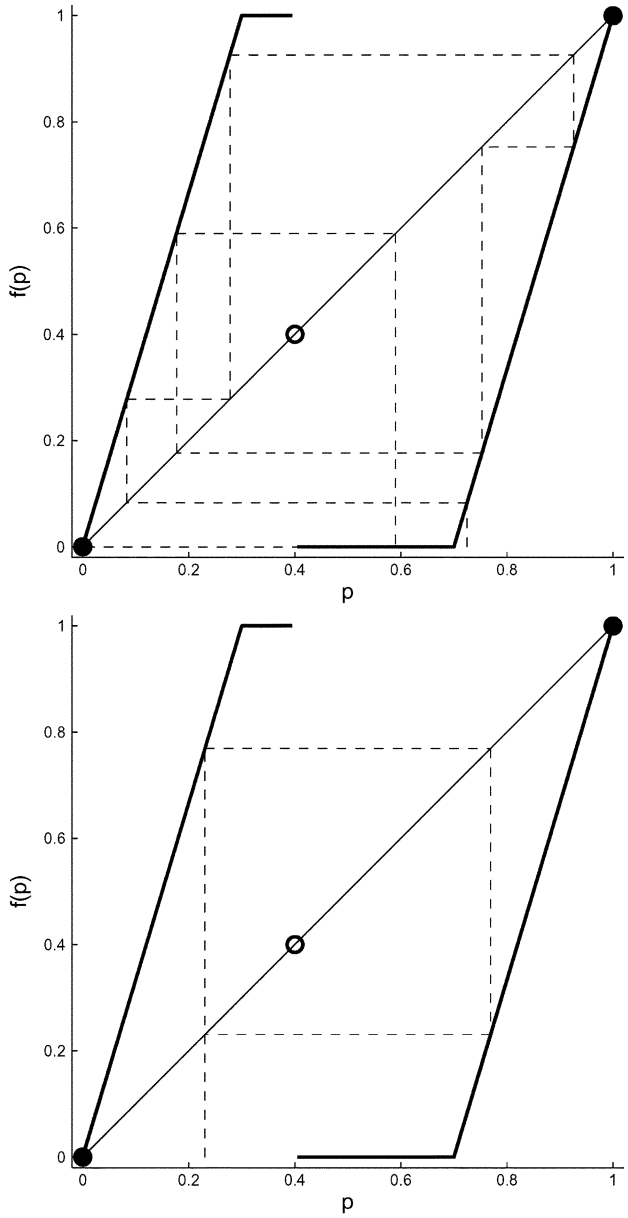

Fig. 8. Operation of (μ, λ) with $p_{\mathrm{EQ}} = 0.4, \gamma = 0.3$. Top: Orbit converges to $p = 0$ ($p_{\mathrm{IC}} = 0.725$). Bottom: Orbit enters unstable two-cycle ($p_{\mathrm{IC}} = 3/13$).

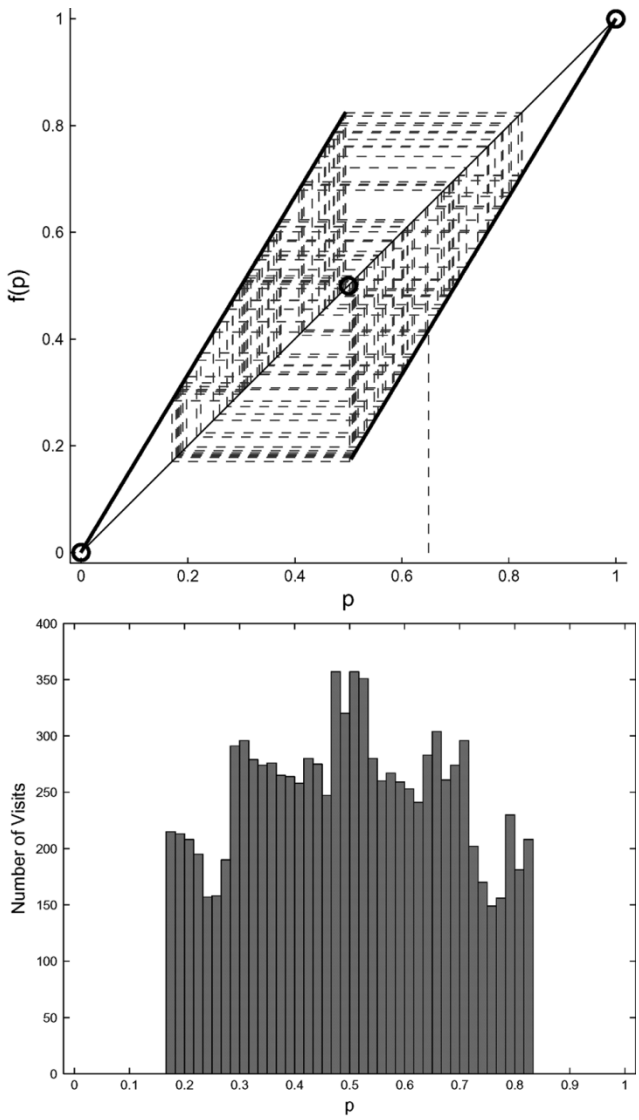

Fig. 9. Operation of (μ, λ) selection with $p_{\mathrm{EQ}} = 0.5, \gamma = 0.6$, and $p_{\mathrm{IC}} = 0.65$. Top: First 100 time-steps of chaotic orbit. Bottom: Histogram of population states visited in chaotic orbit (10^4 time-steps, $\lambda \approx 0.510\,83$).

IX. Linear Ranking

Rank selection is commonly used in *genetic algorithms* [28]–[30]. Ranking proceeds by first sorting the members of the population according to their evaluation scores; the reproductive success of an individual is then a linear or exponential function of its ordinal position, or rank, after sorting. In *linear* ranking, the highest-scoring individual (of a population of size N) receives a rank of N, the next-best $N-1$, and so on; the lowest scoring individual receives rank 1. (When two or more agents have the same evaluation score, then those agents obtain the same rank.) Thus, unlike the other selection methods we examine, linear ranking lacks a parameter with which to vary selection pressure.

Since we have only two strategies in our 1-D framework, linear ranking reduces to giving the higher-scoring strategy a rank of 2 and the other a rank of 1; at polymorphic equilibrium, the strategies obtain the same rank. Plugging these rank values into the standard (i.e., proportional) replicator (6), we obtain our equation for linear ranking (14). The derivatives of our ranking map are given by (15)

$$f(p) = \begin{cases} 2p/(p+1), & \text{if } p < p_{\mathrm{EQ}} \\ p/(2-p), & \text{if } p > p_{\mathrm{EQ}} \\ p_{\mathrm{EQ}}, & \text{if } p = p_{\mathrm{EQ}} \end{cases} \tag{14}$$

$$f'(p) = \begin{cases} 2/(p+1)^2, & \text{if } p < p_{\mathrm{EQ}} \\ 2/(2-p)^2, & \text{if } p > p_{\mathrm{EQ}} \\ \text{undefined}, & \text{if } p = p_{\mathrm{EQ}}. \end{cases} \tag{15}$$

Fig. 10 illustrates the map produced by linear rank selection. As with truncation and (μ, λ), the linear-rank map contains a discontinuity at p_{EQ}; in contrast, the linear-rank map is piecewise smooth, rather than piecewise linear. Also, unlike the selection methods studied above, the linear-rank map yields only a single type of behavior, which is a neutrally stable two-cycle. As soon as an orbit enters the open interval given by (16), the

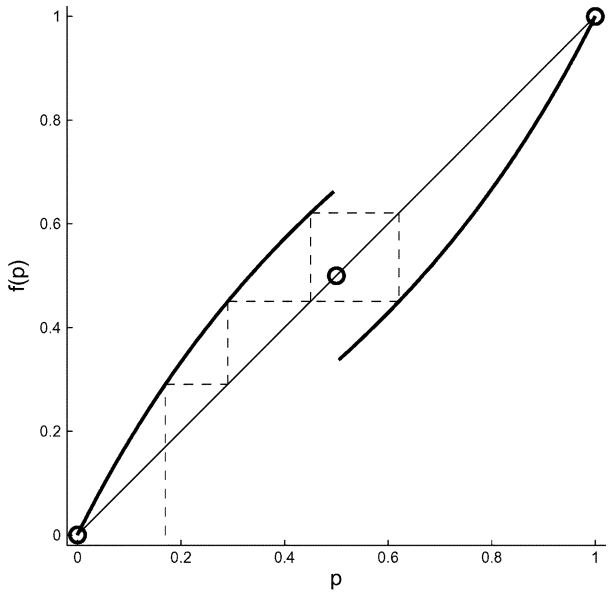

Fig. 10. Map diagram of linear rank selection ($p_{\mathrm{EQ}} = 0.5$). Orbit enters neutrally stable two-cycle ($p_{\mathrm{IC}} = 0.17$).

cycle begins, crossing p_{EQ} with each iteration. The endpoints of the interval (16) belong to $f^{-1}(p_{\mathrm{EQ}})$, that is, the preimage of the linear-rank map at p_{EQ}. Once in a cycle, the slope $f'(p)$ is the inverse of the slope $f'(f(p))$, thus giving a Liapunov exponent of zero for any even number of time-steps. This inverse relationship is shown (for $p < p_{\mathrm{EQ}}$) by (17). Exponential ranking produces qualitatively identical behavior in our framework

$$\left(\frac{p_{\mathrm{EQ}}}{2-p_{\mathrm{EQ}}}, \frac{1+4p_{\mathrm{EQ}}-\sqrt{8p_{\mathrm{EQ}}+1}}{2p_{\mathrm{EQ}}}\right) \tag{16}$$

$$f'(f(p)) = \frac{2}{(2-\frac{2p}{p+1})^2} = \frac{(p+1)^2}{2} = \frac{1}{f'(p)} \quad \text{for } p < p_{\mathrm{EQ}}. \tag{17}$$

X. Best-of-Group Tournament

The first version of tournament selection that we study is described in [11] and [31]; we will term this method "best-of-group" tournament selection. To create a population of N offspring from the current population of N evaluated individuals, we: 1) randomly draw (with replacement) $1 \leq \gamma \leq N$ individuals from the current population; 2) have the fittest member of this subset parent one offspring; and 3) repeat steps 1 and 2 for N times. Selection pressure increases with group size γ. Benefits of this selection method include ease of parallelization and algorithmic efficiency.

Since we only have two strategies in our Hawk–Dove game, the subset produced by step 1 (above) produces one of three outcomes: all Hawks, all Doves, or some blend of the two. If the proportion of Hawks in our population is p, then the probabilities of these three outcomes are $p^\gamma, (1-p)^\gamma$, and $1 - p^\gamma - (1-p)^\gamma$, respectively. Clearly, each all-Hawk subset contributes a Hawk offspring to the next generation; each all-Dove subset contributes a Dove. If $p < p_{\mathrm{EQ}}$, then Hawks outscore Doves and each blended subset contributes a Hawk; similarly, if $p > p_{\mathrm{EQ}}$, then each blended subset contributes a Dove. Thus,

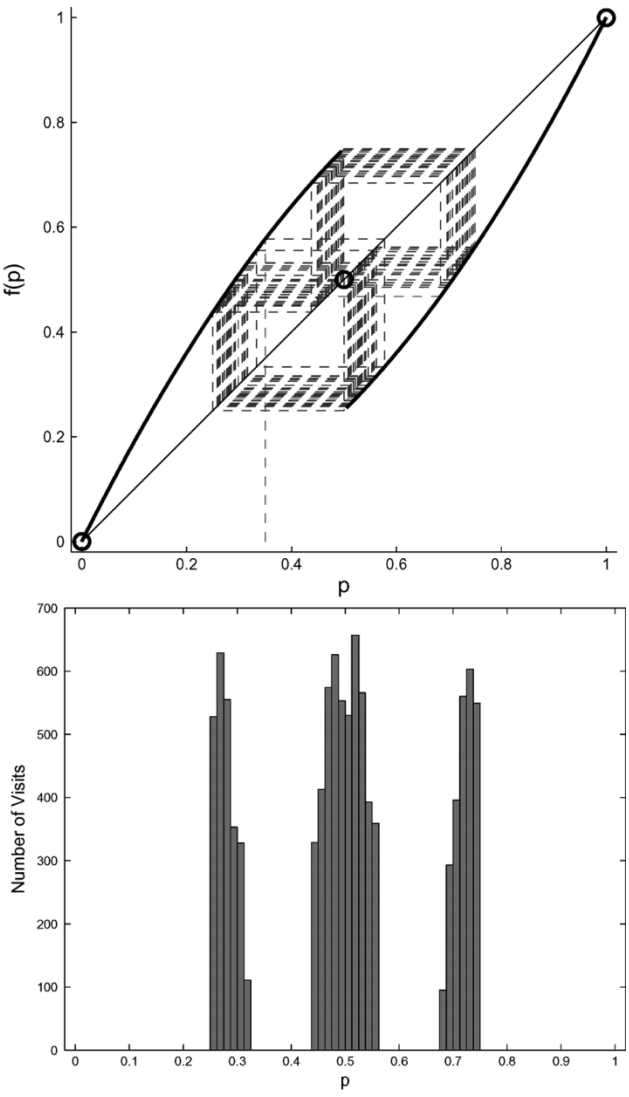

Fig. 11. Map diagram of best-of-group tournament selection ($p_{\mathrm{EQ}} = 0.5, \gamma = 2, p_{\mathrm{IC}} = 0.35$). Top: First 100 time-steps of chaotic orbit. Bottom: Histogram of population states visited in chaotic orbit (10^4 time-steps, $\lambda \approx 0.20979$).

if $p > p_{\mathrm{EQ}}$, then only all-Hawk subsets contribute Hawks, giving us a new population of p^γ Hawks. The map equation for best-of-group tournament selection is given by (18), with corresponding derivatives in (19)

$$f(p) = \begin{cases} 1-(1-p)^\gamma, & \text{if } p < p_{\mathrm{EQ}} \\ p^\gamma, & \text{if } p > p_{\mathrm{EQ}} \\ p_{\mathrm{EQ}}, & \text{if } p = p_{\mathrm{EQ}} \end{cases} \tag{18}$$

$$f'(p) = \begin{cases} \gamma(1-p)^{\gamma-1}, & \text{if } p < p_{\mathrm{EQ}} \\ \gamma p^{\gamma-1}, & \text{if } p > p_{\mathrm{EQ}} \\ \text{undefined} & \text{if } p = p_{\mathrm{EQ}} \end{cases}. \tag{19}$$

The map produced by best-of-group tournament selection has the familiar discontinuity at p_{EQ}, but is otherwise smooth, like linear-ranking. Unlike linear ranking, virtually all initial conditions yield chaotic orbits. Fig. 11 illustrates a chaotic orbit, where $p_{\mathrm{EQ}} = 0.5$ and $\gamma = 2$; note the similarity to Fig. 6 (middle) and (bottom). We can predict a positive Liapunov exponent for this orbit from (19), since it gives $f'(p \neq 0.5) > 1$. As we increase the subset size γ, the map moves away from

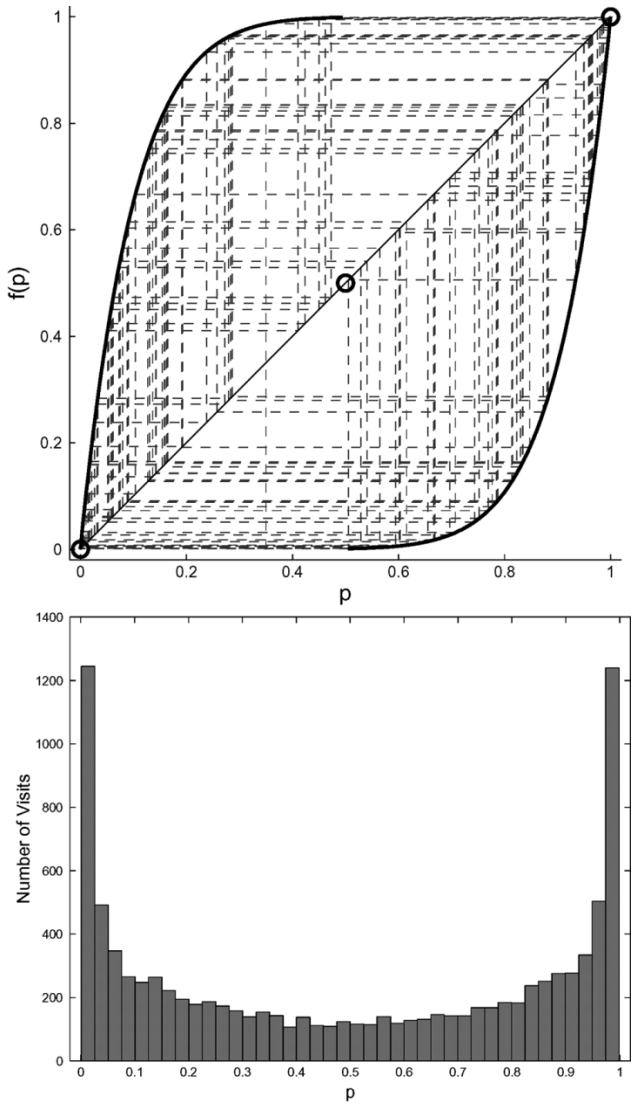

Fig. 12. Map diagram of best-of-group tournament selection ($p_{\mathrm{EQ}} = 0.5, \gamma = 10, p_{\mathrm{IC}} = 0.35$). Top: First 100 time-steps of chaotic orbit. Bottom: Histogram of population states visited in chaotic orbit (10^4 time-steps, $\lambda \approx 0.503\,61$).

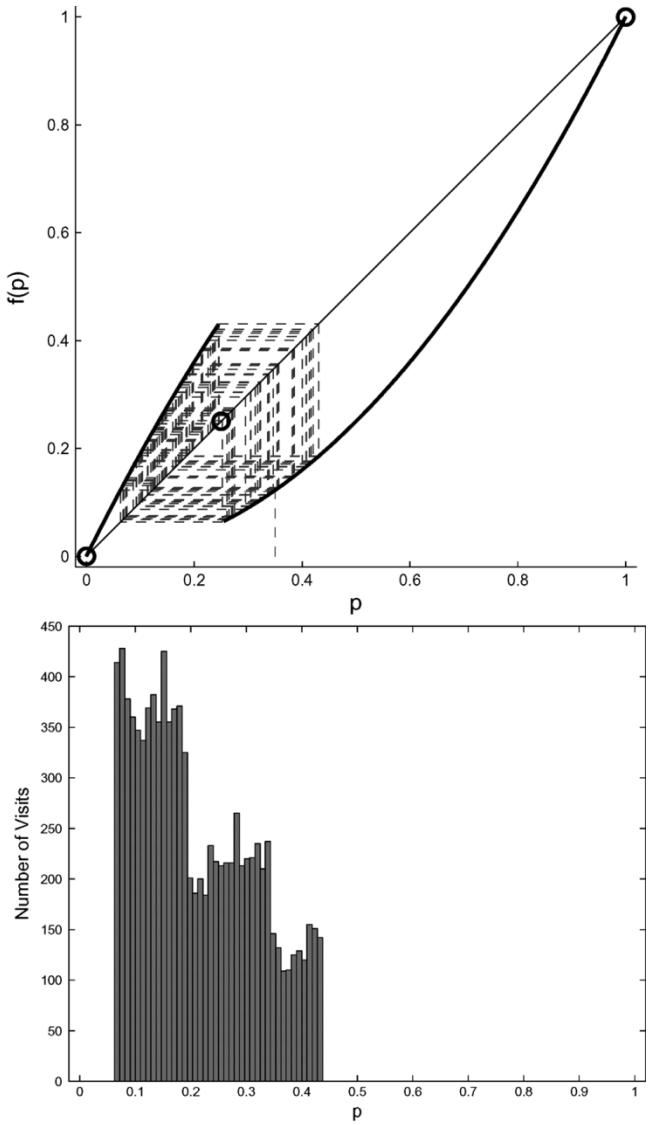

Fig. 13. Map diagram of best-of-group tournament selection ($p_{\mathrm{EQ}} = 0.25, \gamma = 2, p_{\mathrm{IC}} = 0.35$). Top: First 100 time-steps of chaotic orbit. Bottom: Histogram of population states visited in chaotic orbit (10^4 time-steps, $\lambda \approx 0.193\,88$).

the diagonal, causing orbits to spend more time near extrema, where the map is steepest and thereby increasing the Liapunov exponent (Fig. 12).

Orbits that are not chaotic are unstable cycles (or belong to the iterated preimage of the unstable fixed-point p_{EQ}). We find n-cycles by looking for fixed-points of the nth iterate of the map, such that $f^n(p) = p$. For $\gamma = 2$, a two-cycle requires $(2p - p^2)^2 = p$ (for $p < p_{\mathrm{EQ}}$). Solving numerically, we obtain $p_{\mathrm{Cycle}} \approx 0.381\,96$ and $f(p_{\mathrm{Cycle}}) \approx 0.618\,03$. Since a two-cycle must cross p_{EQ} with each iteration, p_{Cycle} will only yield a two-cycle if $0.381\,96 < p_{\mathrm{EQ}} < 0.618\,03$. Higher values of γ contain longer cycles; for example, with $p_{\mathrm{EQ}} = 0.5$ and $\gamma = 10$, unstable cycles of periods three and five exist.

Fig. 13 illustrates another chaotic map, this time with $p_{\mathrm{EQ}} = 0.25$ and $\gamma = 2$. Chaotic orbits (discarding the initial transient) are contained in the open interval $({p_{\mathrm{EQ}}}^2, 2p_{\mathrm{EQ}} - {p_{\mathrm{EQ}}}^2)$. The Liapunov exponent is $\lambda \approx 0.193\,88$, but we notice that $f'(p_{\mathrm{EQ}} < p < 2p_{\mathrm{EQ}} - {p_{\mathrm{EQ}}}^2) < 1$ (while $f'({p_{\mathrm{EQ}}}^2 < p < p_{\mathrm{EQ}}) > 1$). Looking closely at the map, we find that $f(p_{\mathrm{EQ}} < p < 2p_{\mathrm{EQ}} - {p_{\mathrm{EQ}}}^2) < p_{\mathrm{EQ}}$; that is, whenever our chaotic orbit is above p_{EQ}, it will go below p_{EQ} in the very next time-step. A similar statement cannot be made for $f({p_{\mathrm{EQ}}}^2 < p < p_{\mathrm{EQ}})$—we may spend up to three consecutive time-steps below p_{EQ} before crossing it. Thus, we obtain a distribution of population states visited by the orbit that yields a positive Liapunov exponent.

XI. Boltzmann Selection

Inspired by the technique of simulated annealing [32], *Boltzmann selection* is a method whereby selection pressure is slowly increased over evolutionary time to gradually focus search [11], [33], [34]. Selection pressure is modulated by a "temperature" parameter, which is used to scale the rating an agent receives from evaluation; the lower the temperature, the stronger the selection pressure. The scaling equation is $w_{\mathrm{Scaled}} = e^{w/T}$, where w is the agent's unscaled rating and T is the temperature. Agents are selected in proportion to

their scaled ratings. The most effective "cooling schedule" and temperature range are generally problem dependent.

In our discussion, we substitute $\gamma = 1/T$ such that selection pressure increases with γ. The difference equation for our 1-D system using Boltzmann selection is given in (20); the derivative of this map is given in (21). Of all the alternatives to proportional selection that we examine in this paper, only Boltzmann selection is continuous and differentiable throughout the map, as seen in Fig. 14. Nevertheless, Boltzmann selection also exhibits multiple regimes of behavior as γ is varied. (See Yi *et al.* [59] for analysis of this replicator dynamic where γ is fixed at 1.0.) Throughout this section, we use the Hawk–Dove game, where $p_{\mathrm{EQ}} = 7/12$ $(V = 7, C = 12, w_0 = 2.5)$

$$f(p) = \frac{pe^{\gamma(pa-pb+b)}}{pe^{\gamma(pa-pb+d)} - pe^{\gamma(pc-pd+d)} + e^{\gamma(pc-pd+d)}} \tag{20}$$

$$f'(p) = -e^{\gamma(p(a-b+c-d)+b+d)} \times \frac{\gamma(a-b-c+d)(p^2-p)-1}{\left(pe^{\gamma(pa-pb+b)} - pe^{\gamma(pc-pd+d)} + e^{\gamma(pc-pd+d)}\right)^2}. \tag{21}$$

A. Regime I: Point Attractors

For low selection pressures, Fig. 14 (top) shows that the polymorphic Nash-equilibrium attractor at p_{EQ} is intact, and all initial conditions $p_{\mathrm{IC}} \in (0,1)$ converge asymptotically to p_{EQ}. As we increase selection pressure, $f'(p_{\mathrm{EQ}})$ decreases and eventually becomes negative. Dynamical systems theory tells us that, if the slope of a map at a fixed-point is negative but greater than -1, then the fixed-point remains stable, but orbits will show damped oscillation [21]. Such is the case in Fig. 14 (bottom). Oscillation amplitude increases with the steepness of $f'(p_{\mathrm{EQ}})$.

B. Regime II: Limit Cycles

Higher selection pressures make $f'(p_{\mathrm{EQ}}) < -1$ and the fixed-point p_{EQ} loses stability. At $\gamma = 2.0$, where $f'(p_{\mathrm{EQ}}) \approx -1.916\,66$, a stable *limit-cycle* emerges [Fig. 15 (top)]. A stable limit-cycle is a periodic attractor (as opposed to a point attractor, such as p_{EQ} in Fig. 14). In this case, we have a period-two cycle that alternates between $p_{\mathrm{Cyc1}} \approx 0.206\,82$ and $p_{\mathrm{Cyc2}} \approx 0.959\,84$. Because it is an attractor, any initial condition $0 < p_{\mathrm{IC}} < 1$ (provided it is not in the iterated preimage of p_{EQ}) will asymptotically converge onto this same two-cycle, and the cycle will recover if perturbed.

We can easily see the stability of the limit cycle in Fig. 15 (middle); this graph shows the map diagram of the second-iterate $f^2(p)$, which we obtain by composing the Boltzmann map (20) with itself. The map $f^2(p)$ includes the familiar set of fixed-points ($p = 0, p = p_{\mathrm{EQ}}$, and $p = 1$), all of which are unstable. In addition to these three, we find two new fixed-points, which correspond to the alternate values of our limit cycle p_{Cyc1} and p_{Cyc2}. The derivatives of $f^2(p)$ at these new fixed-points have an absolute value less than one, indicating that the fixed-points (and, hence, the cycle) are stable.

We can look for period-three cycles, as well. We simply examine the third-iterate $f^3(p)$, shown in Fig. 15 (bottom). In this case, we find no fixed-points aside from our original three; thus, we know that no three-cycle exists in $f(p)$.

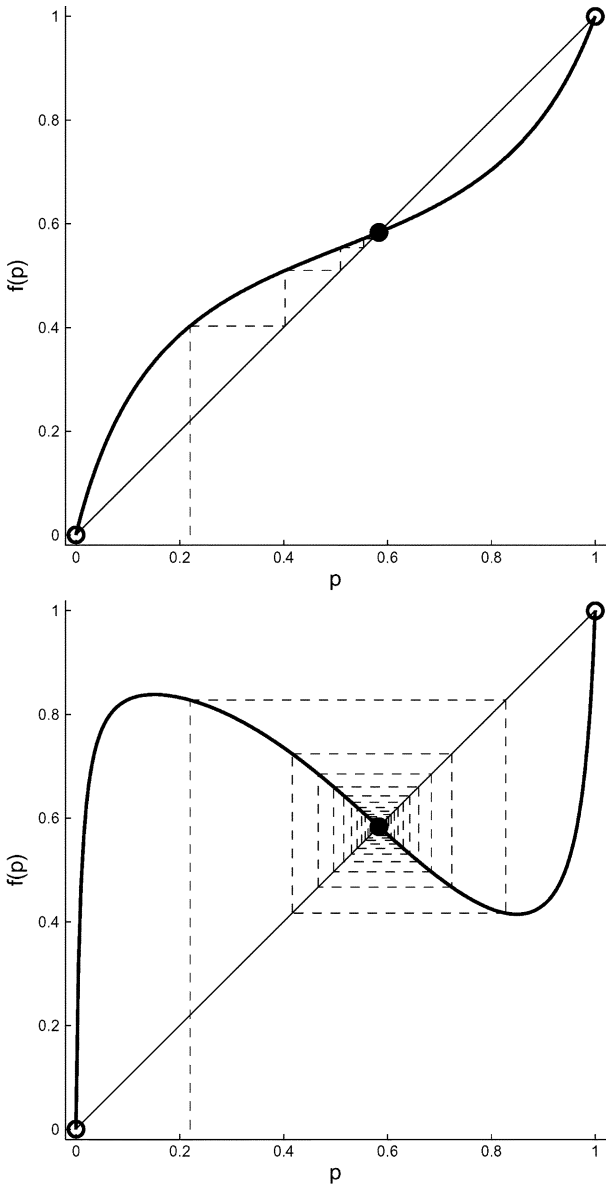

Fig. 14. Map diagrams showing operation of Boltzmann selection ($p_{\mathrm{EQ}} = 7/12, p_{\mathrm{IC}} = 0.22$). Top: Asymptotic convergence to Nash equilibrium attractor is intact ($\gamma = 0.4$). Bottom: Orbit oscillates strongly but still converges to attractor ($\gamma = 1.3$); slope at attractor is $f'(p_{\mathrm{EQ}}) \approx -0.895\,83$.

C. Regime III: Transition to Chaos

Yet higher selection pressure causes our map to become chaotic. Fig. 16 (top) shows one such chaotic orbit when $\gamma = 5.0$; the Liapunov exponent is $\lambda \approx 0.436\,94$. The corresponding histogram of population states visited by the orbit is quite different from those seen above.

The progression of the Boltzmann map from stability to chaos is much more intricate than Figs. 14–16 suggest. A clearer picture emerges from the *orbit diagram* and corresponding graph of Liapunov exponents in Fig. 17. The y axis of the orbit diagram shows the population states that the map visits during its orbit (excluding the initial transient) for a particular value of γ

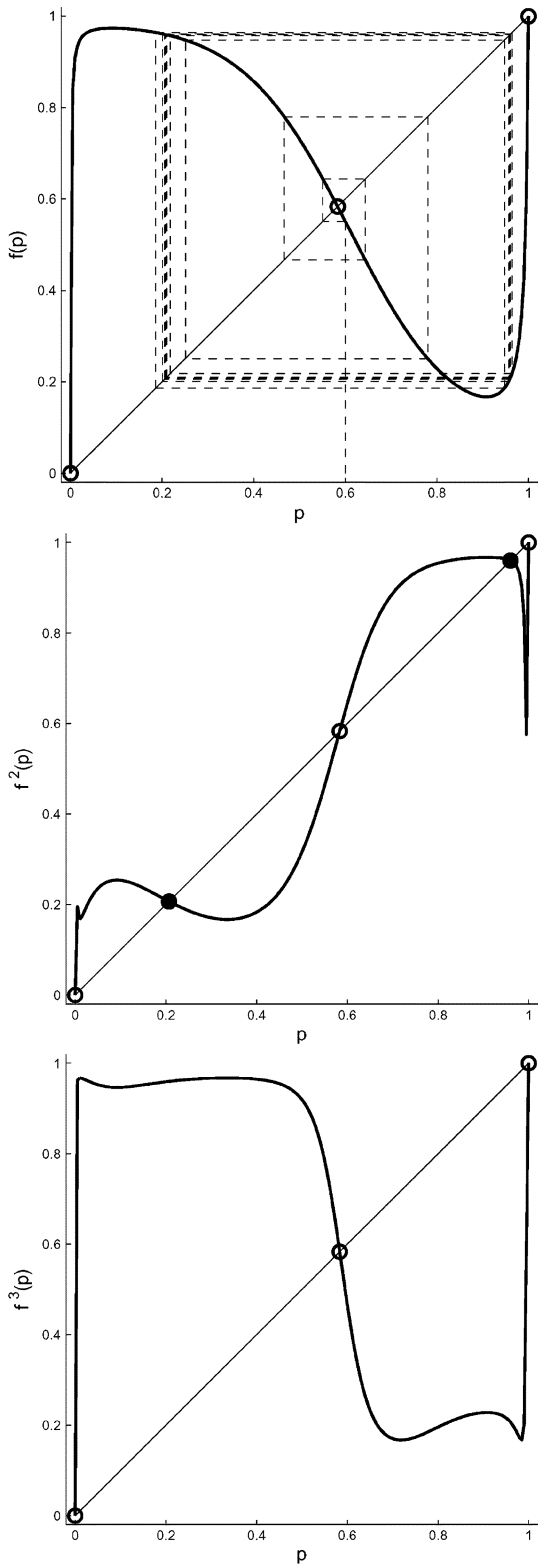

Fig. 15. Map diagrams showing operation of Boltzmann selection on game with $p_{\mathrm{EQ}} = 7/12$. Top: At $\gamma = 2.0$, slope at p_{EQ} becomes less than -1 causing p_{EQ} to become an unstable fixed-point; map now yields stable limit-cycle behavior ($p_{\mathrm{IC}} = 0.6$). Middle: Map diagram produced by second-iterate map $f^2(p)$. Note the two new fixed-points, which correspond to the limit cycle in $f(p)$. Bottom: Map diagram produced by third-iterate map $f^3(p)$. Only the three fixed-points of $f(p)$ are seen; thus, $f(p)$ (top) contains no three-cycles.

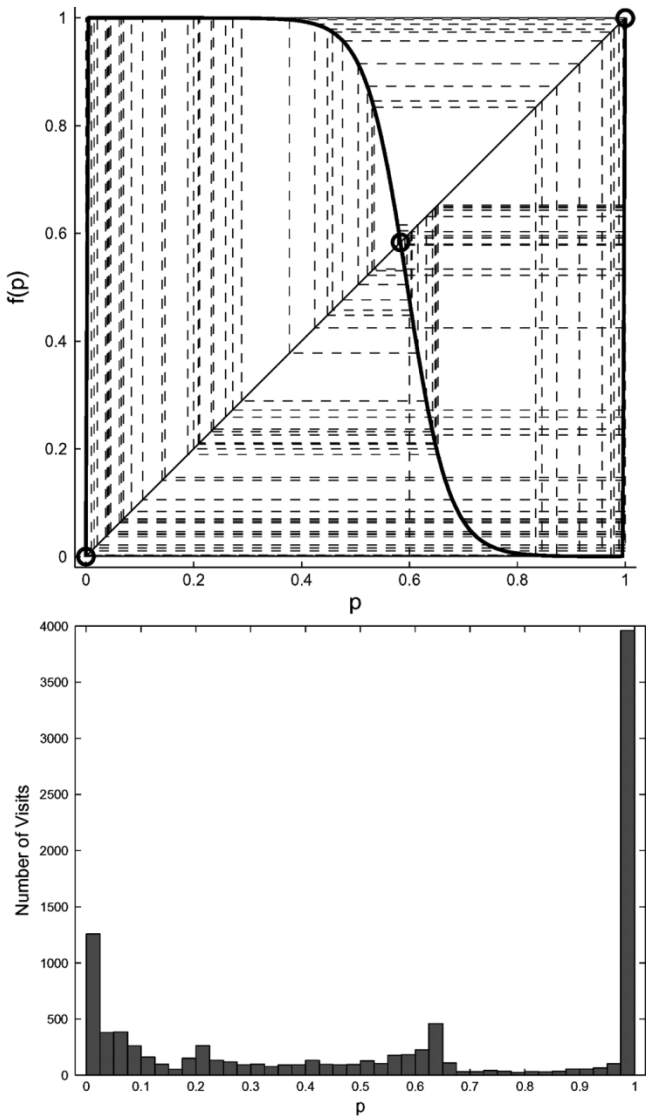

Fig. 16. Map diagram of Boltzmann selection ($p_{\mathrm{EQ}} = 7/12, \gamma = 5.0, p_{\mathrm{IC}} = 0.6$). Top: First 100 time-steps of chaotic orbit. Bottom: Histogram of population states visited by chaotic orbit (10^4 time-steps, $\lambda \approx 0.436\,94$).

(indicated by the x axis). Though the Boltzmann map is not unimodal, we find the classic patterns of bifurcation, period-doubling, and windows of stable oscillation interspersed with chaos that are well-known to occur with unimodal maps (e.g., the logistic map) [21], [35].

We see from the orbit diagram that the fixed-point $p_{\mathrm{EQ}} = 7/12$ bifurcates, in a continuous fashion, into a stable two-cycle at $\gamma \approx 1.36$. That is, the two-cycle appears as a continuous transformation of the polymorphic fixed-point. The next bifurcation, which leads to a period-doubling four-cycle, occurs at $\gamma \approx 2.46$; this process is also continuous. Note that these period-doubling events coincide with the Liapunov exponent approaching zero. As γ exceeds 3.5, the period-doubling process gradually gives way to chaos, and the Liapunov exponent rises above zero. Remarkably, however, further increase of γ brings about a sudden return to stable oscillatory behavior, quite unlike the more gradual process of transition into chaos. These bands of chaos and stable oscillation alternate as γ is increased.

For any unimodal map, if we write down the period of each cyclic region in the order of appearance, then we obtain

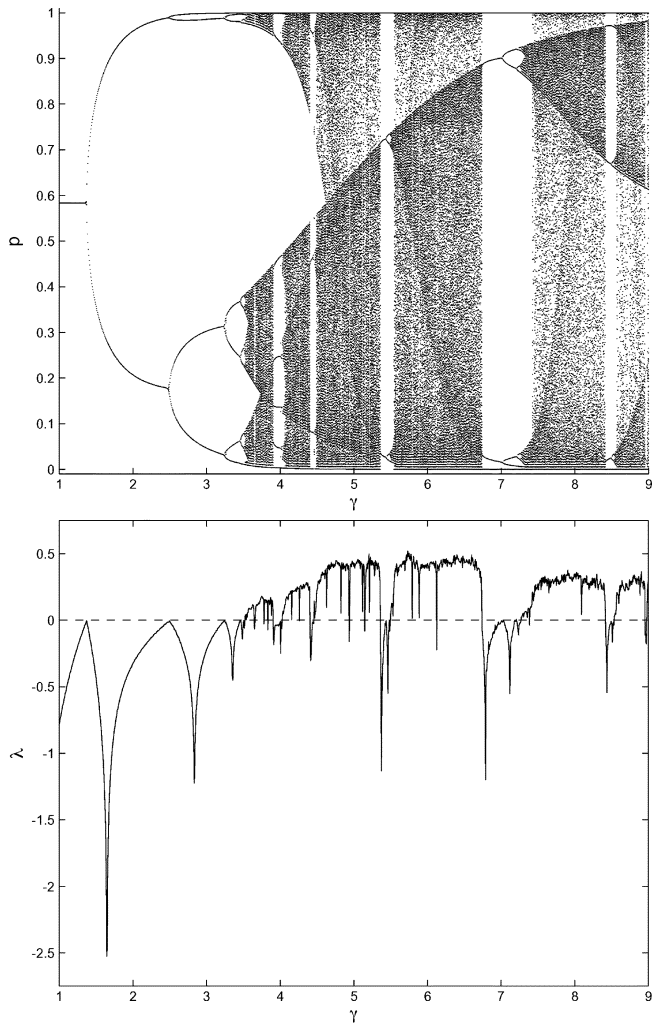

Fig. 17. Orbit diagram (top) and Liapunov exponents (bottom) generated by Boltzmann selection ($p_{\mathrm{EQ}} = 7/12$) over the range $\gamma \in [1, 9]$ with increments to γ of $1/160$.

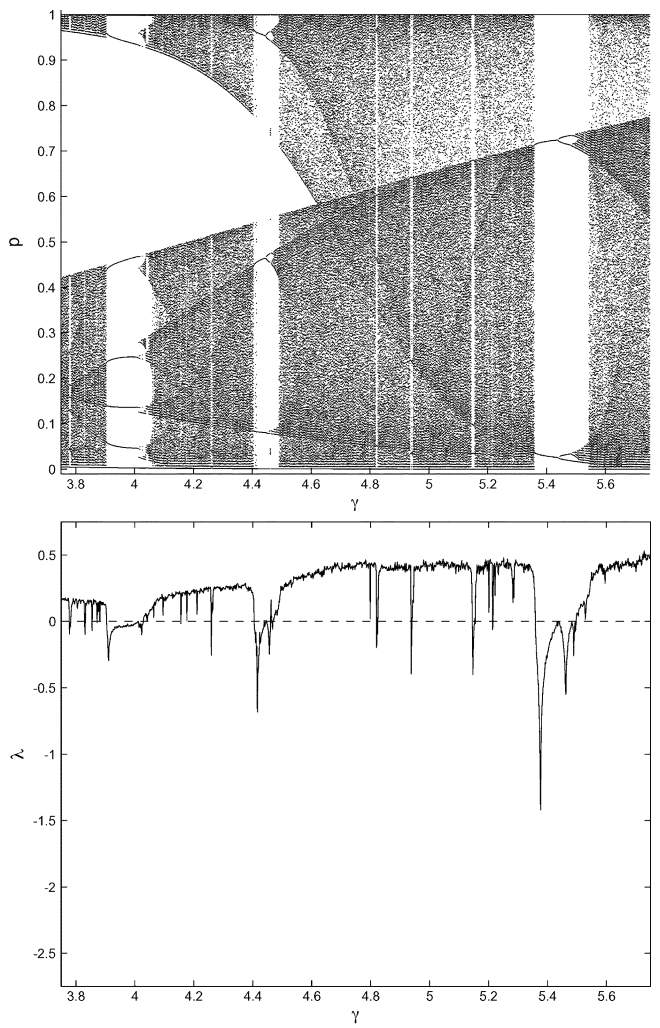

Fig. 18. Orbit diagram (top) and Liapunov exponents (bottom) generated by Boltzmann selection ($p_{\mathrm{EQ}} = 7/12$) over the range $\gamma \in [3.75, 5.75]$ with increments to γ of $1/640$.

precisely the same sequence; this is known as the *universal sequence* [21], [35]. Our Boltzmann map is not unimodal, however, and we find some divergence from the U-sequence. Fig. 18 shows a magnification of the interval $\gamma \in [3.75, 5.75]$. At $\gamma = 3.95$, we appear to have an attracting limit-cycle of period seven, which is different than the six-cycle expected by the U-sequence. Closer inspection reveals this oscillation to actually have a period of ten, with four of the cycle's points very near 1.0. Other cycles include a six-cycle at $\gamma = 4.41$ (which appears in the orbit diagram to be the five-cycle expected by the U-sequence), and a three-cycle at $\gamma = 5.4$ (which corresponds to the U-sequence).

To summarize, we find that Boltzmann selection can produce a variety of behaviors depending upon the strength of the selection pressure. These behaviors include convergence to the polymorphic Nash-equilibrium at p_{EQ}, stable cycles, and chaos. Convergence to p_{EQ} is obtained only for the lowest range of selection pressures. Thus, the analogy between low annealing temperature and high selection pressure is strained—too low a "temperature" actually destabilizes the system.

XII. Sequential-Tournament Selection

The other tournament-style selection method we examine in this paper is often used in evolutionary programming [36]. This version of tournament selection proceeds as follows.

1) Evaluate all individuals in the population.
2) For each individual i, generate an ordered sequence S_i of $\gamma \geq 1$ "opponent" individuals, each randomly drawn (with replacement) from the same population as individual i.
3) If an opponent in S_i has a higher evaluation score than individual i, then let $j \in 1, \ldots, \gamma$ be the index of the earliest such opponent in the sequence S_i; the value $j - 1$ is the *tournament score* earned by individual i. If no such opponent in S_i exists, then the tournament score earned by individual i is γ.
4) Select individuals for reproduction according to their tournament scores. Higher values of γ increase selection pressure. To avoid confusion with the previous tournament method, we will term this approach the "sequential tournament."

Typically, the selection method used in step 4 above is a form of truncation. Given our simple infinite-population framework of two strategies, composing truncation selection with the sequential tournament yields behavior that is identical to truncation selection. Out of curiosity, we instead explore the effects of using proportional selection in step 4. The replicator equation for our approach is given in (22). The corresponding derivatives are given in (23) and (24) for $0 < p < p_{\mathrm{EQ}}$ and $p_{\mathrm{EQ}} < p < 1$, respectively. The map is discontinuous at $p = 0, p = 1$, and p_{EQ}; consequently, the derivative is undefined at these points

$$f(p) = \begin{cases} \frac{\gamma p}{\gamma p + (p-1)(1-p)^{\gamma} - p + 1}, & \text{if } p < p_{\mathrm{EQ}} \\ \frac{p^{\gamma+2} - p^{\gamma+1} - p^2 + p}{(p-1)(p^{\gamma+1} + \gamma p - p - \gamma)}, & \text{if } p > p_{\mathrm{EQ}} \\ p_{\mathrm{EQ}}, & \text{if } p = p_{\mathrm{EQ}} \end{cases} \tag{22}$$

$$f'(p) = \frac{\gamma - (1-p)^{\gamma}(\gamma^2 p + \gamma)}{(\gamma p + (p-1)(1-p)^{\gamma} - p + 1)^2} \tag{23}$$

$$\begin{aligned} f'(p) = {} & (\gamma^2 p^{\gamma+4} - p^{\gamma+3}(4\gamma^2 + \gamma) + p^{\gamma+2}(6\gamma^2 + 3\gamma) \\ & - p^{\gamma+1}(4\gamma^2 + 3\gamma) + p^{\gamma}(\gamma^2 + \gamma) \\ & + \gamma p^3 - 3\gamma p^2 + 3\gamma p - \gamma) \\ & /(p-1)^3/(p^{\gamma+1} + \gamma p - p - \gamma)^2. \end{aligned} \tag{24}$$

To understand the operation of this replicator, let us consider population states $p < p_{\mathrm{EQ}}$. As we know, if $p < p_{\mathrm{EQ}}$, then complete mixing leaves Hawks with a higher cumulative payoff than Doves. Therefore, a Hawk has a cumulative payoff at least as high as any other agent it meets in the sequential tournament, giving the Hawk a tournament score of γ. In contrast, a Dove agent increments its tournament score with each match until it meets a Hawk. Given γ matches and a probability p of meeting a Hawk in each match (and probability $q = 1 - p$ of meeting a Dove), the expected score for a Dove is given by (25). Similar reasoning applies to population states $p > p_{\mathrm{EQ}}$ where Doves outscore Hawks

$$p \sum_{k=1}^{\gamma-1} k q^k + \gamma q^{\gamma} = p \frac{(\gamma-1)q^{\gamma+1} - \gamma q^{\gamma} + q}{(q-1)^2} + \gamma q^{\gamma}. \tag{25}$$

Sequential-tournament selection produces a variety of attracting limit-cycles. For example, Fig. 19 (top) shows an attracting two-cycle for $p_{\mathrm{EQ}} = 0.5$ and $\gamma = 1$. The amplitude of this cycle increases with γ. A period-three attractor is obtained with $p_{\mathrm{EQ}} = 0.05$ and $\gamma = 10$. Fig. 19 (bottom) shows a period-five attractor obtained with $p_{\mathrm{EQ}} = 0.125$ and $\gamma = 10$. These periodic attractors are easily verified by inspecting the second, third, and fifth iterates of the map, respectively.

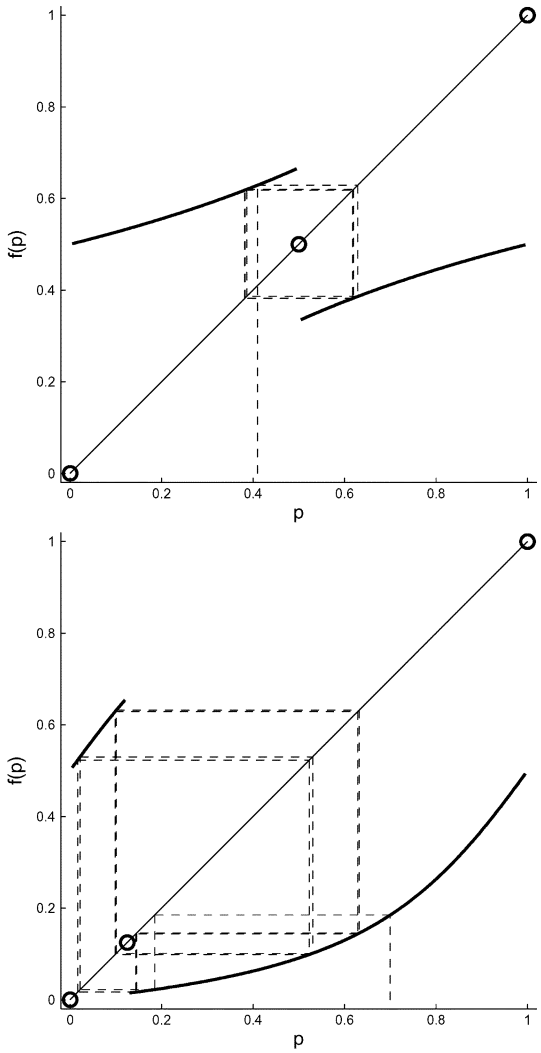

Fig. 19. Attracting limit-cycles produced by sequential-tournament selection. Top: Period-two attractor obtained with $p_{\mathrm{EQ}} = 0.5$ and $\gamma = 1$. Bottom: Period-five attractor obtained with $p_{\mathrm{EQ}} = 0.125$ and $\gamma = 10$.

XIII. Single-Elimination Tournament

Unlike the other selection methods discussed in this paper, the single-elimination tournament [37] is designed specifically for coevolutionary algorithms operating on competitive (e.g., zero-sum) games. In this method, the agents of a population of size N are randomly paired. Each agent that receives the higher game-payoff in pairwise interaction continues to the next round (in case of tie, one agent is chosen at random to win the round); in the second round, the $N/2$ winners of the first round are randomly paired. This process continues until a single, tournament-winning agent remains. The tournament score for each agent corresponds to the highest round the agent won. These tournament scores then form the basis of selection using some other selection method, such as proportional or ranking. Thus, the single-elimination tournament is not only an alternative selection method, but also an alternative to the standard protocol of complete mixing. As an alternative mixing pattern, the single-elimination tournament requires only $O(N \log_2 N)$ matches—a considerable savings if a match is computationally expensive.

In our simple framework, this selection (and mixing) method always converges onto a monomorphic population. The behavior of single-elimination tournament can be understood without deriving the method's equational form. The tournament's pair-wise interactions do not take the population state into account. (In contrast, the other tournament methods operate after agent mixing has occurred, which means that the outcome of a tournament match reflects the state of the population.) Thus, which strategy takes over the population depends exclusively upon the off-diagonal payoffs b and c—payoffs a and d are irrelevant. If $b > c$, then no X-strategist can lose

to a Y-strategist, making the ultimate winner (and most fit individual) an X-strategist *regardless* of the population state p and the equilibrium point p_{EQ}. Similarly, if $b < c$, then a Y-strategist will win. Clearly, this behavior does not correspond to our polymorphic Nash equilibrium.

XIV. Discussion

Our analyses have thus far focused exclusively on the behavior of selection methods around polymorphic equilibria that have a negative-feedback payoff structure. But, the replicator equations we give are completely general and can be applied to any 2×2 symmetric game. Each of the other payoff structures, however, causes selection (regardless of how it maps cumulative payoff to reproductive success) to operate in a positive feedback loop. The behavior of *any* selection method, under positive feedback, will be the following in our simplified model: the strategy with the highest cumulative payoff will both increase in the population *and* maintain a payoff advantage, ultimately taking over the entire population. (See Section IV for details on payoff structures and positive and negative feedback.)

Clearly, real-world coevolutionary algorithms do not adhere to the assumptions of our simplified model. Nevertheless, the dynamics we observe in our model fundamentally underlie those produced by the complications of real-world algorithms (such as genetic variation, finite populations, and noisy evaluation). In particular, the pathologies we observe regarding convergence onto polymorphic Nash equilibria are generally not remedied by the additional complexities of real-world algorithms, but rather are aggravated by them.

As we state in Section I, the conclusions to draw from our results depend very much upon the intent with which the coevolutionary algorithm is used. The algorithm may be used as a search method, a model, or considered purely as a dynamical system. Let us consider the first two possibilities in more depth.

A. Coevolution as Search Method

When used as a search method, or problem solver, the coevolutionary algorithm is most naturally applied to solve games of strategic interaction (for example, [1] and [37]–[42]; many ostensibly static optimization problems can also be reframed as games of strategy, for example, [43]–[45]). The fundamental solution concept for games of strategy is Nash equilibrium [17], [18], which serves as a normative guide: it prescribes the required (or desired) outcome, which in turn drives the search algorithm's design.

Interpreting coevolution as a search method, our results show that many selection methods commonly used in coevolutionary algorithms can locate only monomorphic Nash equilibria (which arise from payoff structures with positive feedback); in the presence of polymorphic Nash arising from negative feedback, these selection methods fail to implement the Nash equilibrium solution concept and prevent the attainment of game-theoretically justifiable results. We emphasize that the lesson learned from our results is not restricted to Nash equilibrium—we may as easily have some other solution concept in mind. Our results, in general, demonstrate that design choices can distort the solution concept that we may otherwise believe to be implemented and operating in our coevolutionary search algorithm. In so doing, we reveal (to our knowledge) a previously unknown class of pathology affecting the use of coevolutionary algorithms as search methods. This paper focuses on design choices that concern selection; other design issues, such as those that concern diversity maintenance, are shown also to impact the solution concept ultimately implemented by a coevolutionary algorithm in [46] and [47].

B. Coevolution as Model

Coevolutionary algorithms can serve as models of other systems. For example, we may construct a coevolutionary algorithm to implement processes that we hypothesize to operate in nature or an economy. In this case, the algorithm's behavior, whatever it may be, is legitimate by definition—at least with respect to the model, if not the system being modeled.

We build models to better understand and predict a system of interest. We wish the model to be no more complicated than necessary to capture the phenomena we wish to study and predict. Thus, we distill the system of interest down to the features that are most salient for our purposes.

Interpreting coevolution as a model, our results emphasize that selection is not a generic process and does not allow the modeler to be indifferent regarding implementation details that concern selection. How we choose to implement selection matters, whether we consider selection to be a focus of our model or not. This point is made concrete in the first of our case studies, below.

Most of the selection methods we test fail to converge onto polymorphic Nash equilibria (though they succeed in converging onto monomorphic Nash). One way to take this result is to question the validity of models such as EGT (which comes to us from the field of evolutionary biology). Specifically, in light of the variety of selection methods that cannot find polymorphisms, EGT may appear to rely too strongly on proportional selection to achieve its results. Before we declare proportional selection to be suspect, however, we must consider the biological plausibility of the alternatives (see [48], for example). We must consider also that proportional selection is not a special case—there exists an infinity of smooth maps that have polymorphic Nash attractors; Boltzmann selection at low pressures is one example (see Appendix C for discussion on a general method to identify smooth maps with stable polymorphisms).

XV. Case Studies

The importance of understanding the dynamics of different selection methods in coevolution is underscored by the following case studies (the first of which stimulated the work we report here). In the light of our results, we believe that the many coevolutionary (and game-theoretic) investigations in the literature that use alternative selection methods may require a second look, especially where a single population is used and the domain under investigation is a variable-sum game.

A. Case One: Evolutionary Game Theory

Our first case study concerns the assumptions made by the EGT framework: an infinitely large population, noiseless payoffs, and complete mixing. Since these assumptions do not hold in the real world, Fogel *et al.* [16], Fogel *et al.* [24], and Fogel and Fogel [49] investigate the effects of finite populations, noisy payoffs, and incomplete mixing upon the equilibria and dynamics that EGT predicts. Using the Hawk–Dove game (with $V < C$) in their simulations, they contrast their empirical results to the polymorphic Nash-equilibrium attractor predicted by EGT.

In addition to the factors concerning population size, payoff noise, and incomplete mixing, all of the simulations in [16], [49], and some of the simulations in [24] deviate from the EGT framework by using truncation selection instead of conventional proportional selection. The results they report from these experiments are consistent with those we discuss in Section VII. In particular, in [49], they observe convergence to all-Hawks when $\gamma = 0.5$ (see our discussion in Section VII-A). In [16], they observe the best correspondence with the polymorphic Nash equilibrium to occur with very low γ (see our Section VII-B). In other experiments in [16], [24], and [49], the authors suspect chaotic behavior. While the dynamics observed in [16], [24], and [49] involve the further complications of finite populations, noisy payoffs, and incomplete mixing, our results and insights from Section VII fundamentally apply. (Readers may note a similarity between [16, Fig. 2] and our orbit diagram in Fig. 17, both of which show an increasing divergence from equilibrium with increasing selection pressure. Nevertheless, the similarities are superficial; for example, our orbit diagrams graph each population state visited in an orbit (after the initial transient), whereas each data point in [16, Fig. 2] is an average population state over an entire run. Instead, the behaviors shown in [16, Fig. 2] are more appropriately contrasted with our predictions for truncation selection over different selection pressures—see Table I for a concise summary. The divergence between [16, Fig. 2] and our analyses is due to the noisy payoffs used in [16]).

Because the experiments in [16], [24], and [49] combine multiple interacting factors, they do not isolate any independent effects of finite populations, noisy payoffs, incomplete mixing, and truncation. We know from our analyses, however, that truncation cannot converge onto a polymorphic equilibrium even when the strong assumptions of EGT are met. Since truncation alone is sufficient to severely distort the dynamics of the EGT framework, the effects of truncation may mask subtler effects involving the other factors of interest.

From the perspective of modeling, we may regard the coevolutionary simulations in [16], [24], and [49] as approximate models of the EGT framework itself; their studies essentially ask how sensitive the EGT framework is to deviations from the ideals provided by EGT's strong assumptions. They are particularly interested in deviations from the conventional EGT framework that move it closer to the natural world. We can therefore justify the use of truncation by asserting that it more closely approximates natural selection (in some system of interest) than proportional selection does.

Certainly, truncation processes exist in nature; for example, predators exert a truncation-like effect on their prey population. But, the specifics of a truncation process are relevant in the determination of biological plausibility. The particular form of truncation selection examined in Section VII (and used in [16], [24], and [49]), along with several other selection methods we examine, operates strictly with ordinal information—a feature that may severely limit its range of plausible applicability. Nevertheless, our results are sufficient to suspect that other, more plausible truncation processes can similarly disrupt the conventional EGT framework. Our results show that coevolutionary algorithm dynamics can be very sensitive to how selection is implemented.

Finally, we note that some of the experiments in [24] do use proportional selection, and statistically significant deviations from Nash equilibrium are observed. Thus, the concerns expressed in [16], [24], and [49] regarding the assumptions of EGT do have merit. Follow-up work on the effects of finite populations provides insight into the mechanism of the distortion [50], [51].

B. Case Two: Iterated Prisoner's Dilemma (IPD)

The large body of research concerning the IPD provides other case studies. The (one-shot) prisoner's dilemma (PD) [52] is a well-known 2×2 symmetric variable-sum game; the payoff structure of the PD game is identical to that of the Hawk–Dove game with $V > C$ in that one of the game's two pure strategies dominates the other (see Sections IV-C and V). (Of course, the PD and Hawk–Dove games are differentiated by other aspects of their payoff structures that do not concern us in this paper.) In the case of the PD game, the Defect strategy always obtains a higher cumulative payoff than Cooperate regardless of how many defectors and cooperators exist in the population. A very different situation can be obtained in the IPD, which is a game where each pairwise interaction between two players consists of a sequence of PD rounds; the number of possible pure strategies of the IPD game grows superexponentially in the number of rounds. A finite collection of IPD strategies (such as might exist in a real-world coevolutionary algorithm's population) is known to be able to have polymorphic equilibrium attractors [53].

In his seminal evolutionary investigation of the IPD, Axelrod [54] uses a selection method that has discontinuities similar to those we find in this paper (this selection method is also used in later work [55]). Agents with cumulative scores within one standard deviation of the population average each receive one offspring; those whose scores are more than one standard deviation above average receive two offspring, and those more than one standard deviation below average receive none. Clearly, this method performs a sorting operation that loses information about payoff convergence. This method is in contrast to proportional selection (i.e., the standard EGT framework) which Axelrod [52, pp. 50–51] uses to analyze his earlier computer-tournament results. (Because the notion of standard deviation is not particularly meaningful for a population composed from just two strategies, we will present our results obtained from Axelrod's selection method elsewhere; the method cannot converge onto polymorphisms.) In other

work, Meuleau and Lattaud [56] use both $(\mu + \lambda)$ and proportional selection in their IPD experiments; they note dramatic differences in their results that are consistent with ours.

XVI. Extensions: Higher-Dimensional Systems

Our investigations above are confined to a single population and the smallest possible game (i.e., two strategies). What generalizations can we make from our results? What happens if we add more strategies or use two populations? We are engaged in new work to answer these questions in detail. Nevertheless, we can say the following about such higher-dimensional systems.

A. Larger Polymorphic Equilibria

First, selection methods such as truncation, $(\mu, \lambda), (\mu + \lambda)$, and ranking cannot converge to a polymorphic equilibrium, regardless of the number of strategies involved in the game (or the polymorphism). The reason is that each of these methods operates by first sorting the population according to their scores. Once the sort is complete, the scores are discarded and only ordinal information is used in the remainder of the selection process. At least two issues arise as a result. To begin, the result of sorting is undefined in the case of ties; how will the agents be ordered if the population is at score-equilibrium? The selection-method definitions we present in (8), (10), and (14) implicitly account for ties by defining the fixed-point at p_{EQ} to exist (albeit as a singularity). Without appropriate modification, standard sorting methods will not maintain this fixed-point. Even if p_{EQ} is accounted for, the ordinal result of sorting throws away information about how similar the agents' scores are. Sorting makes these selection methods insensitive to the process of score convergence, leaving a discontinuity at p_{EQ}. At best, these methods can only provide a cyclic orbit around p_{EQ}. (With an infinite population, the radius of the cycle around p_{EQ} can be made arbitrarily small with truncation, (μ, λ), and $(\mu + \lambda)$ selection, but only by using very weak selection pressure; in a real-world coevolutionary algorithm, finite population size limits the minimal radius. Also, weak selection pressure slows convergence.) Increasing the number of strategies in a polymorphic equilibrium does not smooth the intrinsically discrete result of sorting.

Next, some observations on chaotic behavior. Preliminary experiments show that, as the number of strategies in a polymorphic equilibrium increases, Boltzmann selection requires a stronger selection pressure to be driven to chaos. For example, we can obtain chaotic behavior from a game that has a single polymorphic attractor involving 30 strategies if $\gamma = 350$; we obtain a limit-cycle at $\gamma = 300$, and convergence onto a fixed-point at $\gamma = 200$. (While these selection pressures are well outside of "normal" limits, we must remember that much more reasonable values of γ produce undesirable effects if the polymorphism is small.) Preliminary experiments also indicate that truncation selection becomes prone to chaotic behavior at weaker selection pressures as the number of strategies in a polymorphism increases; cycles and non-Nash fixed-points are also easily observed. Thus, our evidence suggests that the phenomena we present in this paper are not peculiar to our 1-D maps.

B. Two-Population Dynamics

EGT is easily extended to accommodate two populations. The process of complete mixing now requires that each agent in one population interact with each agent in the other population. Agents accumulate payoffs over all interactions. After agent interaction is complete, selection operates on each population independently. Unlike the single-population case, two-population systems do not have polymorphic attractors under the standard (i.e., proportional) replicator dynamics [4]; the polymorphic equilibria we study in this paper remain, but are now unstable. Instead, each population converges to a monomorphic fixed-point. In our Hawk–Dove game, one population will converge to all-Hawks while the other converges to all-Doves. This outcome is also a Nash equilibrium: if one player uses Hawk and the other Dove, then neither has incentive to unilaterally switch strategies.

We can already see that some of our 1-D results apply to this two-dimensional (2-D) system. Truncation selection at $\gamma = 0.5$ provides an easy example. Assume a game with $p_{\mathrm{EQ}} > 0.5$ and proportions of Hawks p_1 and p_2 in populations one and two, respectively. If at time t we have $0.5 \leq p_1, p_2 < p_{\mathrm{EQ}}$ (where p_1 may or may not equal p_2), then both populations will converge to all-Hawks in the next time-step—clearly, not the asymmetric Nash equilibrium we expect. Finally, note that, if the two populations have identical initial conditions, then the dynamical system collapses down to one dimension and acts as if we had only a single population; all of our single-population results apply in this special case.

XVII. Conclusion

We use EGT to investigate selection dynamics in coevolutionary algorithms. We focus on symmetric variable-sum games that have polymorphic Nash equilibria. The behavior of the standard EGT framework, which uses proportional selection, is well studied [4]–[7]; in particular, all point attractors of proportional selection are Nash equilibria, though not all Nash equilibria are attractors. Using the dynamics of proportional selection as a reference, we contrast the behaviors of several commonly used selection methods. The methods we examine are truncation selection, $(\mu, \lambda), (\mu + \lambda)$, Boltzmann selection, linear ranking, and two forms of tournament selection. None of these alternative selection methods (save for Boltzmann) can converge onto any polymorphic Nash equilibrium; further, truncation, (μ, λ), and $(\mu + \lambda)$ selection can have point attractors that are not Nash equilibria. Table I summarizes our results.

Our results are obtained in a simple deterministic framework by the operation of selection alone, without the additional complications introduced by stochastic processes (e.g., genetic variation for search, and random sampling in finite populations). Boltzmann selection is the only method we test that produces a smooth map; the other selection methods produce maps that are piecewise linear or piecewise smooth. In particular, these other selection methods are discontinuous at polymorphic Nash equilibria; further, for most of these discontinuous methods, polymorphic equilibria become fixed points only if we appropriately define the result of sorting in the presence of ties. Even if we take such care, polymorphic equilibria

TABLE I
SUMMARY OF RESULTS

Method	Shape	Discontinuities			γ Range	Constraint	Attractors			Cycles			Chaos
		$p=0$	$p=1$	p_{eq}			$p=0$	$p=1$	p_{eq}	Unstable	Neutral	Attracting	
Proportional	Smooth				n/a	n/a			●				
Truncation	Piece-wise Linear			●	(0, 0.5]	$p_{eq} < \gamma$	●			●			
						$p_{eq} > 1-\gamma$		●		●			
						$\gamma < p_{eq} < 1-2\gamma$					●		
						$2\gamma < p_{eq} < 1-\gamma$					●		
						$1-2\gamma < p_{eq} < 2\gamma$				●	●		●
(μ, λ) and $(\mu+\lambda)$	Piece-wise Linear			●	(0, 1)	$p_{eq} < \gamma \leq 1-\gamma$	●			●			
						$\gamma \leq 1-\gamma < p_{eq}$		●		●			
						$\gamma < p_{eq} < 1-\gamma$	●	●		●			
						$1-\gamma < p_{eq} < \gamma$				●			●
Linear Ranking	Piece-wise Smooth			●	n/a	n/a					●		
Best-of-Group Tournament	Piece-wise Smooth			●	[1, *popsize*]	n/a				●			●
Boltzmann	Smooth				$\mathbb{R}$	Period-doubling route to chaos			●	●		●	●
Sequential Tournament	Piece-wise Smooth	●	●	●	[1, *popsize*]	n/a						●	

are singular points—the discontinuous selection methods fail unconditionally to have polymorphic Nash-equilibrium attractors. Instead, we find point attractors that lack game-theoretic justification (i.e., are not Nash equilibria), cyclic behavior, or chaos. In contrast, Boltzmann selection leaves the polymorphic Nash-equilibrium attractor intact, provided that selection pressure is sufficiently low. For higher selection pressures, attracting limit-cycles and chaos arise.

Our results show that selection methods cannot be moved wholesale from evolutionary to *co*evolutionary frameworks without careful consideration. The selection methods we investigate prevent a coevolutionary algorithm from converging onto polymorphic Nash equilibria, even in a highly idealized setting where we have noiseless evaluation (payoffs), an infinitely large population, and the entire search space is known. Thus, if these selection methods cannot find game-theoretic solutions (specifically polymorphic Nash equilibria) when all the necessary resources are provided, their ability to do so in real-world settings is dubious. (For example, the results of [16], [24], [49] exemplify how additional realism fails to improve the behavior of truncation selection in the Hawk–Dove game.)

A coevolutionary algorithm can be used as a problem solver or as a model of some other dynamical system; our results are meaningful for both cases. With respect to problem solving, all search problems have associated *solution concepts* that define the properties of the solutions we seek (without stating what the solutions are). Though intrinsic to search problems, solution concepts must be *implemented* by search algorithms. The discovery of a solution entails a dynamic adjustment (e.g., search) process. The calibration of the dynamical process to the intended outcome is, therefore, crucial [47].

In this paper, we show that several common selection methods are not "calibrated" to the Nash-equilibrium solution concept. Specifically, we show that polymorphic Nash equi-

libria are problematic. In so doing, we reveal (to our knowledge) a previously unrecognized pathology associated with coevolutionary methods. (We admit the possibility that the action of an ill-calibrated selection method can be counterbalanced by the action of other algorithm components, for example, those that concern variation or diversity maintenance. Nevertheless, such an approach appears far less reliable than simply building upon a properly calibrated selection dynamic.) The pathologies we illustrate are relevant, for example, to many IPD investigations that use coevolutionary methods. All coevolutionary algorithms implement *some* solution concept, either by design or accident. If we construct or use a coevolutionary algorithm without cognizance of the solution concept actually implemented by the algorithm, then we cannot reliably differentiate an unexpected result from algorithm dysfunction.

With respect to modeling, our results highlight the need for fidelity in implementing selection. All of the selection methods we examine share identical attractors in the case of payoff structures with positive feedback (see Section IV). This fact may lead us to believe that selection methods are to some degree interchangeable; indeed, they are all mutually consistent in their mapping from evaluation score to reproductive success (see Section I). Nevertheless, we find that polymorphic equilibria with negative feedback structures cause the selection methods we examine to exhibit very different behaviors—even within our greatly simplified coevolutionary framework. These selection methods are therefore not as interchangeable as we might otherwise believe. Thus, even if the operation of selection is not a focus of our modeling effort, we must exercise care to implement it faithfully. Our results show that coevolutionary dynamics are highly sensitive to the implementation of selection.

Finally, our investigation of selection dynamics reveals unexpected connections with maps studied in the dynamical systems literature [21], [25], [57]. Boltzmann selection follows the well-known period-doubling path to chaos as selection pressure is increased. Truncation, (μ, λ), and $(\mu + \lambda)$ selection (with $p_{\mathrm{EQ}} = 0.5, \gamma = 0.5$) are identical to the binary shift map.

Appendix A
Preimage of a Map

Given a map $f(p)$, let $f^{-1}(p)$ denote the map's *preimage* such that $f^{-1}(p) = \{q : f(q) = p\}$. That is, the preimage of a point p, with respect to the map $f(p)$, is the set of points that go to p when acted upon by the map. Just as we can iterate a map, so too can we iterate the preimage. Let $f^{-k}(p)$ denote the kth iterate of the preimage such that $f^{-k}(p) = \{q : f^k(q) = p\}$. Finally, let $f^{-\infty}(p)$ denote the set of points such that, for each point q in the set, there exists some value of $k > 0$, where $f^k(q) = p$.

Appendix B
Liapunov Exponent

The *Liapunov exponent* λ of a dynamical system is a measure of the system's sensitivity to initial conditions, and is, therefore, an indicator of chaos. For systems that converge to fixed-points or limit-cycles, we obtain $\lambda < 0$. Typically, λ approaches zero as a dynamical system nears a period-doubling bifurcation point. Values $\lambda > 0$ indicate chaos.

As discussed in [21], we approximate the Liapunov exponent with $(1/N)\sum_{t=i}^{N+i-1} \ln(|f'(p_t)|)$. Starting at time $t = i$ (where i is suitably large to allow the initial transient to pass), we sum the natural log of the absolute value of the derivative at each point on the map that our orbit encounters; we then divide by the number of time-steps to get an average. For the Liapunov values we report in this paper, we use $i = 5000$ and $N = 10^4$; the histograms of population states visited in chaotic orbits do not include the initial transient of 5000 time-steps.

Appendix C
Determining Stability in Higher Dimensions

This appendix reviews a standard test that can be applied to any differentiable, multidimensional map M to determine its stability properties [58]. This test is based on the Hartman–Grobman theorem, which allows us to treat a system as if it were linear in the vicinity of the fixed-point. By doing so, we can apply the simple stability tests of linear maps to the fixed points.

Given a symmetric game of n strategies, our selection function is a map M with $m = n - 1$ dimensions. Let $\mathbf{p}_{\mathrm{Fix}} = \{p_1, \ldots, p_m\}$ be a fixed-point of our map, such that $M(\mathbf{p}_{\mathrm{Fix}}) = \mathbf{p}_{\mathrm{Fix}}$. To calculate the stability of this fixed-point, we first linearize the system by calculating its first derivative at $\mathbf{p}_{\mathrm{Fix}}$. Because our map has more than one dimension, we need to calculate its *Jacobian matrix*

$$\partial M(\mathbf{p}_{\mathrm{Fix}}) = \begin{pmatrix} \frac{\partial M_{p_1}}{\partial p_1} & \cdots & \frac{\partial M_{p_1}}{\partial p_m} \\ \vdots & \ddots & \vdots \\ \frac{\partial M_{p_m}}{\partial p_1} & \cdots & \frac{\partial M_{p_m}}{\partial p_m} \end{pmatrix} \tag{26}$$

where $(\partial M_{p_i})/(\partial p_j)$ is the partial derivative of map variable p_i with respect to variable p_j.

The test for convergence is to check whether the (potentially complex) eigenvalues λ of the Jacobian fall within the interior of the unit-circle on the real-complex plane; that is, we check that the real and complex parts of λ have absolute values less than one. If all eigenvalues fall within the unit circle, then the fixed point is stable. If one of the eigenvalues falls outside, then the fixed point is unstable.

To gain some intuition about this test, consider a 1-D map, such as proportional selection in Fig. 2. This map contains a stable fixed-point, such that if we iterate the map from a point near the fixed-point then we will converge onto the fixed-point. Why do points in the neighborhood of the fixed-point converge onto it? Take an initial point at an offset from the fixed point, $p_0 = p_{\mathrm{Fix}} + \mu$, close enough to the fixed point that we can treat the map as linear. The linearization of the map in the region of p_{Fix} allows us to approximate the map as $M(p) \approx p_{\mathrm{Fix}} + \lambda(p - p_{\mathrm{Fix}})$, where λ corresponds to the slope of the map at the fixed point. An iteration of the map has the effect of multiplying the offset by the slope $p_1 = M(p_0) \approx p_{\mathrm{Fix}} + \lambda(p_{\mathrm{Fix}} + \mu - p_{\mathrm{Fix}}) = p_{\mathrm{Fix}} + \lambda\mu$. If $|\lambda| < 1$, then application of the map will cause p_1 to be closer to p_{Fix} than p_0. Therefore, multiple iterations of the map cause the offset to be multiplied repeatedly by the slope,

$p_n = p_{\mathrm{Fix}} + \lambda^n \mu$, and (if $|\lambda| < 1$) bring it closer and closer to the fixed-point.

Thus, the test for a 1-D map is whether the absolute value of the derivative at the fixed point is less than one. The multidimensional test is based on the same convergence properties as the 1-D case. In a way, taking the eigenvalues of the Jacobian matrix is equivalent to breaking the multidimensional system down into constituent 1-D systems, where each eigenvalue represents the rate of change (derivative) of each 1-D degree of freedom of the multidimensional system. These eigenvalues can be complex numbers, rather than reals. Nevertheless, the test stays the same in the sense that we still test whether multiplying an offset by an eigenvalue will shrink the offset. Thus, in the multidimensional case, the test becomes whether all the eigenvalues' magnitudes are less than one.

ACKNOWLEDGMENT

The authors would like to thank L. Altenberg, A. Bucci, and D. Fogel for extensive discussions, the reviewers and editors for many helpful comments, and members of the DEMO Laboratory.

REFERENCES

[1] C. D. Rosin and R. Belew, "New methods for competitive coevolution," *Evol. Comput.*, vol. 5, no. 1, pp. 1–29, 1997.

[2] J. Cartlidge and S. Bullock, "Learning lessons from the common cold: How reducing parasite virulence improves coevolutionary optimization," in *Proc. Congr. Evol. Comput.*, D. B. Fogel, M. A. El-Sharkawi, and X. Yao, Eds, 2002, pp. 1420–1425.

[3] J. M. Smith, *Evolution and the Theory of Games*. Cambridge, U.K.: Cambridge Univ. Press, 1982.

[4] J. Hofbauer and K. Sigmund, *Evolutionary Games and Population Dynamics*. Cambridge, U.K.: Cambridge Univ. Press, 1998.

[5] J. Weibull, *Evolutionary Game Theory*. Cambridge, MA: MIT Press, 1995.

[6] L. Samuelson, *Evolutionary Games and Equilibrium Selection*. Cambridge, MA: MIT Press, 1997.

[7] D. Fudenberg and D. K. Levine, *The Theory of Learning in Games*. Cambridge, MA: MIT Press, 1998.

[8] D. E. Goldberg and K. Deb, "A comparative analysis of selection schemes used in genetic algorithms," in *Foundations of Genetic Algorithms (FOGA 1)*, G. J. Rawlins, Ed., 1991, pp. 69–93.

[9] P. J. Hancock, "An empirical comparison of selection methods in evolutionary algorithms," in *Lecture Notes in Computer Science*, T. C. Fogarty, Ed. New York: Springer-Verlag, 1994, vol. 865, Proc. Evol. Comput. (AISB Workshop), pp. 80–94.

[10] T. Blickle and L. Thiele, "A comparison of selection schemes used in evolutionary algorithms," *Evol. Comput.*, vol. 4, no. 4, pp. 361–394, 1996.

[11] M. Mitchell, *An Introduction to Genetic Algorithms*. Cambridge, MA: MIT Press, 1996.

[12] E. Cantú-Paz, "Order statistics and selection methods of evolutionary algorithms," *Inf. Process. Lett.*, vol. 82, no. 1, pp. 15–22, 2002.

[13] S. G. Ficici, O. Melnik, and J. B. Pollack, "A game-theoretic investigation of selection methods used in evolutionary algorithms," in *Proc. 2000 Congr. Evol. Comput.*, A. Zalzala *et al.*, Eds., 2000, pp. 880–887.

[14] J. M. Smith and G. R. Price, "The logic of animal conflict," *Nature*, vol. 246, pp. 15–18, 1973.

[15] M. A. Nowak, A. Sasaki, C. Taylor, and D. Fudenberg, "Emergence of cooperation and evolutionary stability in finite populations," *Nature*, vol. 428, pp. 646–650, Apr. 2004.

[16] G. B. Fogel, P. C. Andrews, and D. B. Fogel, "On the instability of evolutionary stable strategies in small populations," *Ecological Modeling*, vol. 109, pp. 283–294, 1998.

[17] J. Nash, "Noncooperative games," *Ann. Math.*, ser. 2nd, vol. 54, no. 2, pp. 286–295, 1951.

[18] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA: MIT Press, 1998.

[19] D. Bishop and C. Cannings, "A generalized war of attrition," *J. Theor. Biol.*, vol. 70, pp. 85–124, 1978.

[20] R. Dawkins, *The Selfish Gene*. London, U.K.: Oxford Univ. Press, 1989.

[21] S. H. Strogatz, *Nonlinear Dynamics and Chaos*. Reading, MA: Addison-Wesley, 1994.

[22] B. Skyrms, "Chaos and the explanatory significance of equilibrium: Strange attractors in evolutionary game dynamics," in *Proc. Biennial Meeting Philos. Sci. Assoc.*, vol. 2, 1992, pp. 374–394.

[23] D. B. Fogel, "An overview of evolutionary programming," in *Evolutionary Algorithms*. ser. IMA Volumes in Mathematics and its Applications, L. D. Davis, K. De Jong, M. D. Vose, and L. D. Whitley, Eds. Berlin, Germany: Springer-Verlag, 1997, pp. 89–109.

[24] D. B. Fogel, G. B. Fogel, and P. C. Andrews, "On the instability of evolutionary stable states," *BioSystems*, vol. 44, pp. 135–152, 1997.

[25] H.-O. Peitgen, H. Jürgens, and D. Saupe, *Chaos and Fractals: New Frontiers of Science*. New York: Springer-Verlag, 1992.

[26] T. Bäck, "Evolution strategies: An alternative evolutionary algorithm," in *Lecture Notes in Computer Science*, J.-M. Alliot, E. Lutton, E. Ronald, M. Schoenauer, and D. Snyers, Eds., 1996, vol. 1063, Proc. Eur. Conf. Artif. Evol., pp. 3–20.

[27] H.-P. Schwefel, *Evolution and Optimum Seeking*. New York: Wiley, 1995.

[28] J. Grefenstette, "Rank-based selection," in *Handbook of Evolutionary Computation*, T. Bäck, D. B. Fogel, and Z. Michalewicz, Eds. Bristol, U.K.: Inst. Physics, 1997, ch. C2.4.

[29] J. E. Baker, "Adaptive selection methods for genetic algorithms," in *Proc. 1st Int. Conf. Genetic Algorithms*, J. J. Grefenstette, Ed., 1988, pp. 101–111.

[30] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, MA: Addison-Wesley, 1989.

[31] T. Blickle, "Tournament selection," in *Handbook of Evolutionary Computation*, T. Bäck, D. B. Fogel, and Z. Michalewicz, Eds. Bristol, U.K.: Inst. Physics, 1997, ch. C2.3.

[32] E. Aarts and J. Korst, *Simulated Annealing and Boltzmann Machines: A Stochastic Approach to Combinatorial Optimization and Neural Computing*. New York: Wiley, 1989.

[33] M. de la Maza and B. Tidor, "An analysis of selection procedures with particular attention paid to proportional and Boltzmann selection," in *Proc. 5th Int. Conf. Genetic Algorithms*, S. Forrest, Ed., 1993, pp. 124–131.

[34] S. W. Mahfoud, "Boltzmann selection," in *Handbook of Evolutionary Computation*, T. Bäck, D. B. Fogel, and Z. Michalewicz, Eds. Bristol, U.K.: Inst. Physics, 1997, ch. C2.5.

[35] N. Metropolis, M. Stein, and P. Stein, "On finite limit sets for transformations on the unit interval," *J. Combinatorial Theory (A)*, vol. 15, pp. 25–44, 1973.

[36] D. B. Fogel, "Other selection methods," in *Handbook of Evolutionary Computation*, T. Bäck, D. B. Fogel, and Z. Michalewicz, Eds. Bristol, U.K.: Inst. Physics, 1997, ch. C2.6.

[37] P. J. Angeline and J. B. Pollack, "Competitive environments evolve better solutions for complex tasks," in *Proc. 5th Int. Conf. Genetic Algorithms*, S. Forrest, Ed., 1993, pp. 264–270.

[38] C. W. Reynolds, "Competition, coevolution and the game of tag," in *Artificial Life IV*, R. A. Brooks and P. Maes, Eds. Cambridge, MA: MIT Press, 1994, pp. 59–69.

[39] S. Nolfi and D. Floreano, "Co-evolving predator and prey robots: Do 'arm races' arise in artificial evolution?," *Artificial Life*, vol. 4, no. 4, pp. 311–335, 1998.

[40] J. B. Pollack and A. D. Blair, "Co-evolution in the successful learning of backgammon strategy," *Machine Learning*, vol. 32, no. 3, pp. 225–240, 1998.

[41] K. Chellapilla and D. B. Fogel, "Evolving neural networks to play checkers without expert knowledge," *IEEE Trans. Neural Netw.*, vol. 10, no. 6, pp. 1382–1391, 1999.

[42] D. B. Fogel, T. Hays, S. Hahn, and J. Quon, "A self-learning evolutionary chess program," *Proc. IEEE*, vol. 92, no. 12, pp. 1947–1954, 2004.

[43] D. Hillis, "Co-evolving parasites improves simulated evolution as an optimization procedure," *Physica D*, vol. 42, pp. 228–234, 1990.

[44] L. Pagie and P. Hogeweg, "Evolutionary consequences of coevolving targets," *Evol. Comput.*, vol. 5, no. 4, pp. 401–418, 1997.

[45] H. Juillé and J. B. Pollack *et al.*, "Coevolving the 'ideal' trainer: Application to the discovery of cellular automata rules," in *Proc. 3rd Annu. Genetic Program. Conf.*, J. R. Koza *et al.*, Eds., 1998, pp. 519–527.
[46] S. G. Ficici and J. B. Pollack, "Game theory and the simple coevolutionary algorithm: Some results on fitness sharing," in *Proc. Genetic Evol. Comput. Conf. Workshop Program*, R. Heckendorn, Ed., 2001, pp. 2–7.
[47] S. G. Ficici, "Solution Concepts in Coevolutionary Algorithms," Ph.D. dissertation, Brandeis Univ., Waltham, MA, May 2004.
[48] D. B. Fogel and H.-G. Beyer, "Do evolutionary processes minimize expected losses?," *J. Theor. Biol.*, vol. 207, pp. 117–123, 2000.
[49] D. B. Fogel and G. B. Fogel, "Evolutionary stable strategies are not always stable under evolutionary dynamics," in *Evolutionary Programming IV*, J. McDonnell, R. Reynolds, and D. B. Fogel, Eds. Cambridge, MA: MIT Press, 1995, pp. 565–577.
[50] S. G. Ficici and J. B. Pollack *et al.*, "Effects of finite populations on evolutionary stable strategies," in *Proc. 2000 Genetic Evol. Comput. Conf.*, L. D. Whitley *et al.*, Eds., 2000, pp. 927–934.
[51] ——, "Finite-population dynamics in games with polymorphic attractors," Comput. Sci. Dept., Brandeis Univ., Waltham, MA, Tech. Rep. CS-05-257, 2005.
[52] R. Axelrod, *The Evolution of Cooperation*. New York: Basic Books, 1984.
[53] M. A. Nowak and K. Sigmund, "Game-dynamical aspects of the prisoner's dilemma," *Appl. Math. Comput.*, vol. 30, pp. 191–213, 1989.
[54] R. Axelrod, "The evolution of strategies in the iterated prisoner's dilemma," in *Genetic Algorithms and Simulated Annealing*, L. Davis, Ed. San Mateo, CA: Morgan Kaufmann, 1987, pp. 32–41.
[55] R. L. Riolo, M. D. Cohen, and R. Axelrod, "Evolution of cooperation without reciprocity," *Nature*, vol. 414, pp. 441–443, 2001.
[56] N. Meuleau and C. Lattaud, "The artificial evolution of cooperation," in *Lecture Notes in Computer Science*, J.-M. Alliot, E. Lutton, E. Ronald, M. Schoenauer, and D. Snyers, Eds. New York: Springer-Verlag, 1996, vol. 1063, Proc. Eur. Conf. Artif. Evol., pp. 159–180.
[57] A. Katok and B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems*. Cambridge, U.K.: Cambridge Univ. Press, 1995.
[58] R. W. Easton, *Geometric Methods for Discrete Dynamical Systems*. London, U.K.: Oxford Univ. Press, 1998.
[59] T. Yi, Y. Qi-Sen, J. Zhi-Gang, and W. Zu-Wang, "Evolutionarily stable strategy, stable state, periodic cycle and chaos in a simple discrete time two-phenotype model," *J. Theor. Biol.*, vol. 188, pp. 21–27, 1997.

Sevan G. Ficici received the B.M. degree in music from the Cleveland Institute of Music, Cleveland, OH, in 1988, the M.A. degree in music theory from the Eastman School of Music, Rochester, NY, in 1990, and the Ph.D. degree in computer science from Brandeis University, Waltham, MA, in 2004.

He is currently a Postdoctoral Fellow in computer science at the Artificial Intelligence Research Group, Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA. His research interests include evolutionary dynamics, game theory, multiagent systems, machine learning, robotics, and computer analysis of music.

Ofer Melnik received the B.S. degree in computer science from the Israeli Open University, Tel-Aviv, in 1995 and the Ph.D. degree in computational neuroscience from Brandeis University, Waltham, MA, in 2001, where he was a Sloan Fellow.

He is currently a Postdoctoral Fellow at DIMACS, Rutgers University, Piscataway, NJ. He has publications on machine learning, black-box model interpretation, coevolutionary optimization dynamics, and recurrent neural dynamical systems, including multiple journal papers. His current research interests include the above and machine learning problems related to fusion, representations, and high-dimensional interpretation.

Dr. Melnik received the Best Student Paper Award at the International Joint Conference on Neural Networks in 1998.

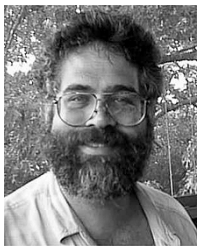

Jordan B. Pollack is Professor of Computer Science and Complex Systems at Brandeis University and Director of the DEMO Laboratory, Waltham, MA. He has broad interests and publications across artificial intelligence, artificial life, neural and evolutionary computing, robotics, complex and dynamic systems, educational technology, and intellectual property law. He has been involved in startup companies including Abuzz, Affinnova, Nannon Technology, and Thinmail.

Dr. Pollack made front page news worldwide in 2000, and was named one of MIT's Technology Review ''TR 10'' in 2001 for his laboratory work on automatically designed and manufactured robots.

Errata

Equation (16) should read:

$$\left(\frac{p_{\mathrm{EQ}}}{2-p_{\mathrm{EQ}}}, \frac{2p_{\mathrm{EQ}}}{1+p_{\mathrm{EQ}}}\right)$$