A Game-Theoretic Investigation of Selection Methods Used in Evolutionary Algorithms

Sevan G. Ficici
sevan@cs.brandeis.edu

Ofer Melnik
melnik@cs.brandeis.edu

Jordan B. Pollack
pollack@cs.brandeis.edu

DEMO Lab—Computer Science Department
Volen Center for Complex Systems, Brandeis University
Waltham Massachusetts USA
www.demo.cs.brandeis.edu

Abstract-

The replicator equation used in evolutionary game theory (EGT) assumes that strategies reproduce in direct proportion to their payoffs; this is akin to the use of fitness-proportionate selection in an evolutionary algorithm (EA). In this paper, we investigate how various other selection methods commonly used in EAs can affect the discrete-time dynamics of EGT. In particular, we show that the existence of evolutionary stable strategies (ESS) is sensitive to the selection method used. Rather than maintain the dynamics and equilibria of EGT, the selection methods we test impose a fixed-point dynamic virtually unrelated to the payoffs of the game matrix, give limit cycles, or induce chaos. These results are significant to the field of evolutionary computation because EGT can be understood as a *coevolutionary algorithm* operating under ideal conditions: an infinite population, noiseless payoffs, and complete knowledge of the phenotype space. Thus, certain selection methods, which may operate effectively in simple evolution, are pathological in an ideal-world coevolutionary algorithm, and therefore dubious under real-world conditions.

1 Introduction

A central achievement of evolutionary game theory (EGT) was the introduction of a method by which agents can play "optimal" strategies in the absence of rationality [Maynard-Smith, 1982]. Through a process of Darwinian selection, a population of agents can evolve to an *evolutionary stable strategy* (ESS), which is a Nash equilibrium with an additional stability criterion. While the dynamics and equilibria in EGT are influenced by the payoff matrix, the system's general dynamical properties are determined by the *replicator equation*. The canonical replicator used in EGT is a difference (or differential) equation that selects agents to reproduce offspring in direct proportion to fitness.

The dynamics of various replicator equations is a topic of intense study, particularly with respect to the game-theoretic equilibria [Weibull, 1995, Samuelson, 1997]. Further, ties between EGT, quantitative genetics, and animal behavior have launched a host of more biologically inspired investigations into the dynamics of replicator systems [Hofbauer and Sigmund, 1998, Dugatkin and Reeve, 1998]. Evolutionary game theory has also lead to a proliferation of various non-evolutionary alternatives to agent rationality that include social learning methods such as imitation [Fudenberg and Levine, 1998]. In all these studies, an understanding of the dynamics of "replication"—the engine of learning—is key.

Yet, the selection methods (replicators) used in evolutionary algorithms (EA) have escaped careful game-theoretic scrutiny. This paper investigates the properties of four common EA selection methods from a game-theoretic and dynamical-systems perspective: truncation, (μ, λ)-ES, linear ranking, and Boltzmann selection. Selection methods used in EAs are reviewed in [Goldberg and Deb, 1991, Hancock, 1994, Mitchell, 1996]. But, how do these methods compare to the canonical fitness-proportionate selection when used in EGT? Do they exhibit similar dynamics and promote the same fixed-points and attractors?

Rather than maintain the dynamics and equilibria of EGT, the selection methods we test impose a fixed-point dynamic virtually unrelated to the payoffs of the game matrix, give limit cycles, or induce chaos. Only Boltzmann selection (and only at low selection pressures) is faithful to EGT dynamics. These results transcend evolutionary game theory, however. Because our analysis is game-theoretic, we are implicitly considering the effects of these selection mechanisms in an evolutionary environment where strategy *interaction* is taken into account—otherwise known as a *coevolutionary algorithm*. We conclude that certain selection methods, while they may be effective in ordinary evolutionary algorithms, are likely to be inappropriate for coevolution.

We begin with a brief introduction to evolutionary game theory and the *Hawk-Dove* game. The next two sections outline our dynamical systems approach, where we first look at the dynamical features of the Hawk-Dove game, and then consider the role of the selection method (the replicator). We then consider each of the four selection methods, in turn. Discussions of how our results generalize to larger games, and what the results mean for coevolutionary algorithms follow. We then review a general methodology used in dynamical systems to determine the stability characteristics of an arbitrary differentiable replicator function. We finish with concluding remarks.

2 Evolutionary Game Theory

This section introduces evolutionary game theory and the *Hawk-Dove* game [Maynard-Smith, 1982]. As explained above, the key contribution of EGT is the notion that Darwinian selection can replace the need for agent rationality. Let us consider a concrete example. The (symmetric) payoff matrix for the Hawk-Dove game, G, is shown in Equation 1. We imagine an infinitely large population of agents, each playing one of the game's two *pure strategies*, 'H' or 'D'. The state of the population can be represented by a column vector, p, that represents the proportion with which each strategy appears in the population; the elements of the vector thus sum to 1.0. The *fitness* of each strategy (and hence of all agents that play that strategy) is a weighted sum of payoffs in G, the weights being determined by the composition of the population. The fitness vector, f, is computed by matrix multiplication, as in Equation 2. To the result of the multiplication, we add the constant, w_0, such that all fitness values are positive. Once the fitness of each strategy is known, the next generation of the population, p_{t+1}, can be created by applying Darwinian selection to the current population, p_t. This process is accomplished by the difference equation in Equation 3. According to this equation, each strategy reproduces in direct proportion to its fitness—evolutionary game theory assumes *fitness-proportionate selection*.

$$G = \begin{array}{c|cc} & \mathrm{H} & \mathrm{D} \\ \hline \mathrm{H} & -25 & 50 \\ \mathrm{D} & 0 & 15 \end{array} \tag{1}$$

$$f = G * p + w_0 \tag{2}$$

$$p_{t+1} = p_t \times \frac{f}{p_t \bullet f} \tag{3}$$

where $w_0 = 26$ is a constant added to fitnesses such that they are all greater than zero, '$\times$' is element-wise multiplication, and '$\bullet$' is inner product. The lower term in Equation 3 is for normalization.

3 The Dynamics of the Hawk-Dove Game

What is the behavior of Equation 3? Let us consider it from a dynamical systems perspective. Dynamical systems theory deals with the behavior of functions with recursive properties, such as difference and differential equations. Its tools are designed to identify and classify various behaviors common to different systems. These behaviors include many classes of fixed-points, bifurcations, and orbits, as well as chaos.

Though the Hawk-Dove game has two strategies, it is actually a one-dimensional system: given the proportion of one strategy, we can deduce the proportion of the other: $p^{doves} = 1 - p^{hawks}$. For this reason, we can visualize the game's dynamical properties with a *map diagram*, shown in Figure 1. The curve in Figure 1 depicts the function that maps the proportion of Hawks at time t, to the proportion of Hawks in the next generation, $t+1$. Where the curve of a one-dimensional map crosses the diagonal line, there exists a *fixed point*; that is, $p_{t+1} = \mathrm{Map}(p_t)$. This implies that all strategies that participate in a fixed point must receive the same fitness, otherwise their proportions would have to change. Mathematically, we can denote a fixed point as $f = c\overrightarrow{1}$, for any constant c, which implies a vector of strategy proportions $p = (G + w_0)^{-1} c\overrightarrow{1}$. The Hawk-Dove game has three fixed points: 1) the population is composed of all Hawks, 2) the population is all Doves, and 3) where $7/12^{ths}$ of the population is Hawk and $5/12^{ths}$ is Dove.

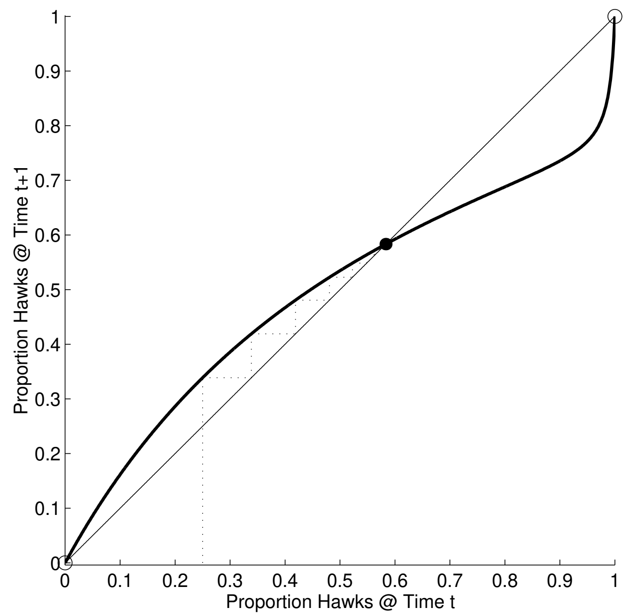

Figure 1: Hawk-Dove game under EGT replicator dynamics.

In a one-dimensional map, the slope of the curve at a fixed point determines the *stability* characteristics of the fixed point. If the slope has an absolute value less than one, then the fixed point is *stable*; if the absolute value is greater than one, then the fixed point is *unstable*. A dynamical system that is started with an initial condition, p_0, in the neighborhood (basin of attraction) of a stable fixed-point, p^{fix}, will converge to the fixed-point. Mathematically we can write this as: if $\left\| p_0 - p^{\mathrm{fix}} \right\| < \varepsilon$ (for a specific ε), then $\lim_{t \to \infty} p_t = p^{\mathrm{fix}}$. In the Hawk-Dove game, the two fixed points where the population is composed entirely of a single strategy are unstable. The third fixed point is stable, and is an *evolutionary stable strategy* (ESS). All evolutionary stable strategies are stable fixed points. If the population at an ESS is composed of more than a single pure strategy, as it is in the Hawk-Dove game, then the ESS is *polymorphic*.

Figure 1 can be used to create a *cobweb* diagram, which depicts the trajectory, or *orbit*, of an arbitrary initial condition. To find the orbit of an initial condition, p_0, we begin on the x-axis at p_0 and draw a vertical line to the curve; then we alternately draw a horizontal line to the diagonal and a vertical line to the curve to trace the evolution of the population for each subsequent time step. In the Hawk-Dove game, under Equation 3, we find that any initial condition other than $p_0 = 0$ and $p_0 = 1.0$ will converge onto the ESS.

4 The Role of Selection

We should understand that the dynamics of evolutionary game theory are not determined solely by the payoff matrix of the game, G. Indeed, Equation 3 can be rewritten as the composition of two constituent functions, which describes a more general difference equation:

$$p_{t+1} = M(p_t) = S(F(p_t), p_t) \tag{4}$$

The column vector p represents strategy proportions in the population, as before. The map M is the composition of two sub-functions. The function $F(p_t)$ is the *fitness function*, which calculates each strategy's fitness, given the current state of the population, p_t. In evolutionary game theory, the fitness function is Equation 2, and this is where the payoff matrix exerts its influence. But, there is a second function involved. The function $S(f, p_t)$ is the *selection function*, which takes vectors of strategy fitnesses, f, and corresponding proportions, p_t, of the current generation, and returns the state of the population for the next generation. In evolutionary game theory, S performs fitness-proportionate selection.

This paper examines what happens to EGT dynamics when the selection mechanism, S, is changed: when are the stable-state dynamics maintained and when are they perturbed? Any S that maintains the stable fixed points of Equation 3 must obey the identity $p^{\text{fix}} = S(c\overrightarrow{1}, p^{\text{fix}})$, as well as show convergence properties around p^{fix}. Failure to meet these two necessary and sufficient conditions will inevitably disrupt the ESS properties of the system. The alternatives to fitness-proportionate selection that we investigate are four selection methods commonly found in evolutionary algorithms: truncation, (μ, λ)-ES, linear rank, and Boltzmann selection.

5 Truncation Selection

Truncation selection is used primarily in the branch of evolutionary computation known as *evolutionary programming* [Fogel, 1997]. Given a population of size n and selection pressure k, truncation selection operates by first sorting the population according to fitness and then replacing the worst k percent of the population with variations of the best k percent. For example, for a population size of $n = 200$ and selection pressure of $k = 25$, variations of the best 50 individuals in the population will be copied over the worst 50 individuals. Thus, selection pressure must be in the range $0 \leq k \leq 50$, with higher values of k giving higher selection pressure. Since evolutionary game theory excludes variational operators, we simply replace the worst $k\%$ with exact copies of the best $k\%$.

Truncation selection is easily implemented for the infinite population assumed by evolutionary game theory. Since all agents that play the same strategy will receive the same fitness, we only need to sort the strategies and note the proportions with which each strategy appears in the population. Given a game of m strategies, we represent the state of the population with a vector, p, of size m, where p^i represents the proportion of strategy i in the population. Sorting these strategies yields a new vector, q, where q^{s_i} is the proportion of strategy, s_i, with sorted rank i. To select the agents to remove, we construct a vector w (representing the worst $k\%$) that indicates the proportion of each strategy we are to remove from the population. Another vector, b (representing the best $k\%$), indicates the proportion of each strategy we are to add to the population. The new proportions, q', are $q' = q - w + b$. Figure 2 gives an example where q is comprised of three strategies, each used by 1/3 of the population. Selection pressure is at the maximum of $k = 50\%$. We see that the worst 50% of the population is composed of all agents playing strategy s_3 and one half of those playing s_2. Similarly, the best half of the population is all those agents that play s_1 and half of those that play s_2. In the new population, therefore, strategy s_1 composes 2/3 of the population, s_2 1/3 of the population, and strategy s_3 is eliminated.

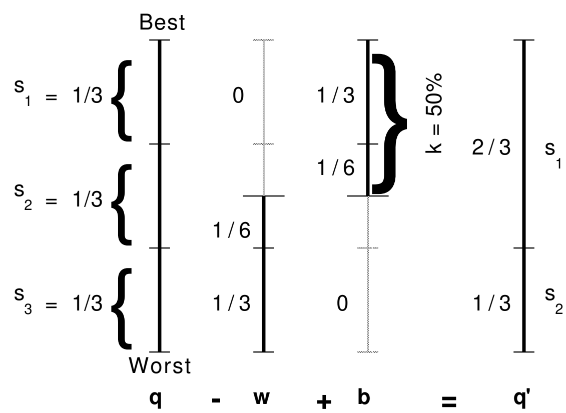

Figure 2: Truncation selection on an infinite population.

How does truncation selection behave when the population contains more than one strategy and all strategies present in the population receive the same fitness? Such a condition implies the existence of a fixed point, such as a polymorphic ESS. Since the result of sorting is ill-defined in the case of ties, so too is the operation of truncation selection. Unless special precautions are taken to deal with ties in fitness, truncation selection is unable to maintain arbitrary fixed points.

How are the dynamics of the Hawk-Dove game affected when we replace fitness-proportionate selection with truncation selection? We find that truncation selection exhibits three modes of behavior over the range of selection pressure.

For selection pressures $42\% \leq k \leq 50\%$, truncation selection causes most initial conditions to converge to all

Hawks, though some initial conditions lead to cyclic behavior. The map diagram produced by truncation with $k = 50$ is shown in Figure 3 (top). We know that at the ESS proportion of 7/12, both strategies receive the same payoff. If the proportion of Hawks is below 7/12, then Hawks receive higher fitness than Doves; and, if the proportion of Hawks is above 7/12, then Doves receive higher fitness than Hawks. If fitness-proportionate selection were used, this would produce a simple feedback mechanism that would keep the population at the ESS proportion.

Truncation, however, breaks this feedback mechanism. To understand why, let us consider a population state where the proportion of Hawks is $\frac{1}{2} \leq p^{\mathrm{H}} < \frac{7}{12}$. Because p^{H} is below the ESS, the Hawks will receive higher fitness than the Doves. But, the Hawks also comprise at least half of the population. Thus, the best 50% of agents in the population can only be playing the Hawk strategy, and the next generation will be 100% Hawks. At $p^{\mathrm{H}} = \frac{7}{12}$, we have already noted that truncation cannot maintain the fitness equilibrium. Thus, the orbit of almost every initial condition that includes both strategies, $0 < p_0^{\mathrm{H}} < 1$, will eventually fall into the critical interval $\frac{1}{2} \leq p^{\mathrm{H}} < \frac{7}{12}$, and therefore converge to the attractor of all Hawks. An example of an initial condition that leads to a cycle of period two is $p^{\mathrm{H}} = \frac{2}{3}$ at $k = 50\%$.

For selection pressures $31\% \leq k \leq 41\%$, truncation causes most initial conditions to have chaotic orbits, while some give cycles. Figure 3 (middle) shows one such chaotic orbit where $k = 36$. An indicator of chaos is the *Liapunov exponent*, λ, which measures sensitivity to initial condition. The Liapunov exponent is negative for fixed-point dynamics, approaches zero for limit cycles, and is greater than zero for chaos. It is normally calculated by measuring the derivative of the map at each point in an orbit. Because the truncation map has discontinuities, it is not differentiable. Nevertheless, it is piece-wise linear, and the lack of smoothness is negligible. Therefore, we use the slope of the line segment. This yields a Liapunov exponent of $\lambda = 0.69$.

If the selection pressure is in the range $0 < k \leq 30\%$, all initial conditions $0 < p_0^{\mathrm{H}} < 1$ result in cyclic behavior. A sample orbit is shown in Figure 3 (bottom) with $k = 15$. We should take note of the map's discontinuity at the ESS proportion: for $p^{\mathrm{H}} < \frac{7}{12}$, the map is above the diagonal; for $p^{\mathrm{H}} > \frac{7}{12}$, the map is below. This means that for a cycle to occur, the ESS proportion must be crossed. Thus, all cycles go around the ESS. Nevertheless, because the upper and lower segments of the map are parallel to the diagonal, the exact location of the cycle is determined by where the orbit first enters the cycle inducing region of the map. As selection pressure is decreased, the cycles exhibit tighter orbits around the ESS proportion.

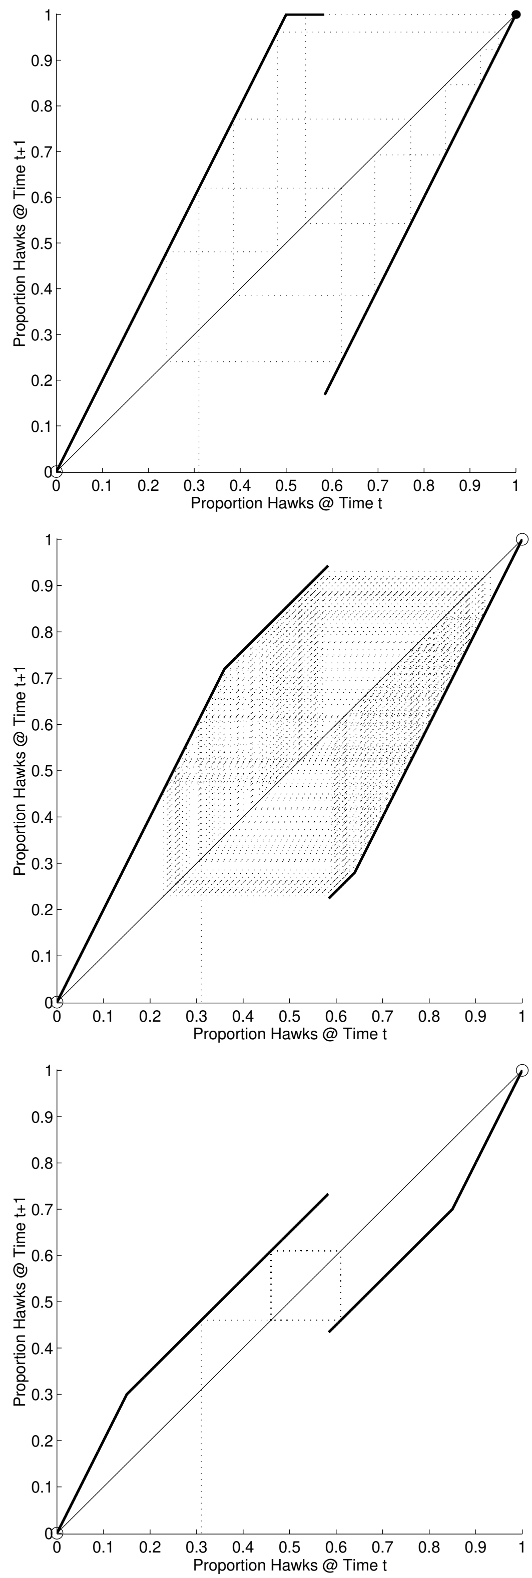

Figure 3: Map diagram of truncation selection with selection pressure at 50% (top), 36% (middle), and 15% (bottom).

6 (μ, λ)-ES Selection

The (μ, λ)-ES selection method is used in the branch of evolutionary computing known as *evolution strategies*

[Bäck, 1995]. Given a population of λ offspring, the best μ offspring are chosen to parent the next generation. Normally, variational operators are applied during reproduction, but we omit variation here, as we did with truncation selection. The (μ, λ)-ES selection method is similar to truncation selection, but more drastic. In (μ, λ)-ES selection, the best $\frac{\mu}{\lambda}\%$ expands to replace the entire population rather than just the worst $\frac{\mu}{\lambda}\%$.

The implementation of (μ, λ)-ES selection for infinite populations is identical to that of truncation selection, except that now vector w is discarded and vector b is normalized to create the new population. The fraction μ/λ determines the selection pressure—the actual values of μ and λ are unimportant for an infinite population. Thus, for (μ, λ)-ES selection, the selection pressure can be in the range $0 < \frac{\mu}{\lambda} \leq 1.0$, where lower values indicate higher selection pressure.

For the Hawk-Dove game, (μ, λ)-ES selection also has three regimes of behavior, none of which are able to maintain arbitrary fixed points and all of which include some initial conditions that lead to cycles. For the range $.59 \leq \frac{\mu}{\lambda} < 1.0$, the behavior is usually chaotic. Figure 4 (top) shows an example orbit where $\frac{\mu}{\lambda} = 0.6$. The measured Liapunov exponent is 0.52. In the range $.42 \leq \frac{\mu}{\lambda} \leq .58$, the system converges to all Hawks in a manner very similar to truncation selection. The range $0 < \frac{\mu}{\lambda} \leq .41$ introduces the additional possibility of converging onto all Doves. This is shown in Figure 4 (bottom), where $\frac{\mu}{\lambda} = 0.3$. Because the map of (μ, λ)-ES selection is not differentiable, the stability properties of the all-Hawk and all-Dove fixed points are unusual—they are attractors, but they are not locally stable. The same is true for the all-Hawk attractor seen in truncation selection.

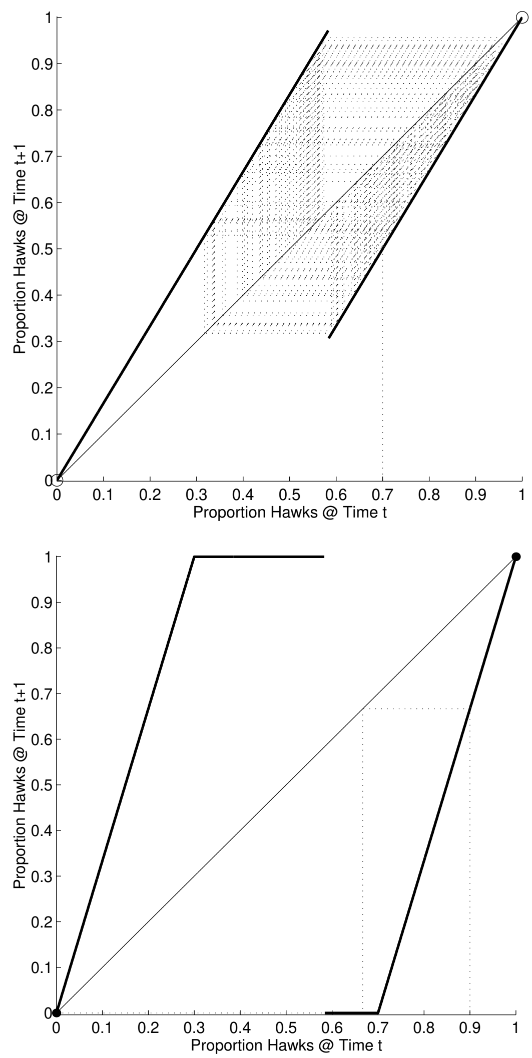

Figure 4: Map diagram of (μ, λ)-ES selection with $\frac{\mu}{\lambda} = 0.6$ (top) and 0.3 (below).

7 Linear Rank Selection

Ranking is a method commonly used in *genetic algorithms* [Mitchell, 1996]. Agents are sorted according to fitness, and then assigned *new* fitness values according to their rank. In *linear* ranking, the new assigned fitness values change linearly with rank. Fitness-proportionate roulette-wheel selection is then applied using the new fitness values. This method is useful in ordinary genetic algorithms because, as a population converges, the differences between agents' fitnesses can become too small for the roulette wheel to resolve, given the population size. Ranking ensures that small the differences between fitness levels are expanded. Ranking also attenuates very large differences between fitness levels, which helps prevent premature convergence.

No special modifications are needed to implement linear ranking for our infinite population: we rank strategies according to fitness, and then assign new values. Because the Hawk-Dove game has only two strategies, the strategy with lower fitness is assigned a new fitness of one, and the other strategy gets a new fitness of two. After normalization, the new fitnesses are either [1/3; 2/3], [2/3; 1/3], or, if special care is taken to handle ties (at the ESS), [.5; .5]. We then proceed with proportional selection on these new fitnesses.

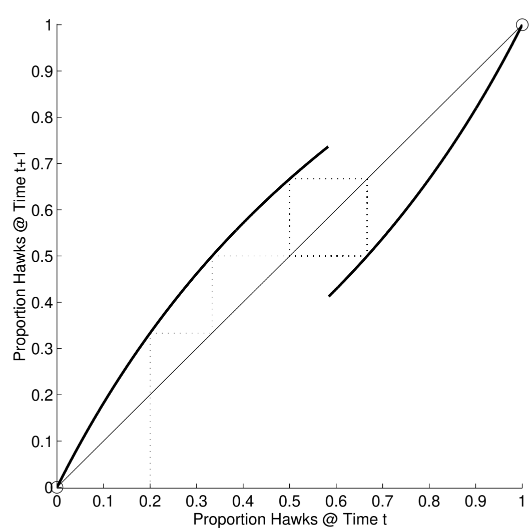

Figure 5: Map diagram of linear rank selection method.

As Figure 5 shows, linear ranking produces cycles around the ESS. In fact, this is the only behavior that linear ranking can exhibit. The rates of change for the two strategies are exactly inverted as the ESS proportion is crossed, thus ensuring a simple period-two cycle. Because ranking maps all possible fitness proportions to a single proportion, rates of change never approach equality; thus, ranking cannot have attractive fixed points. This is visible in the map diagram. The map never crosses the diagonal—there is only a discontinuity at the ESS proportion. Another version of rank-based selection assigns new fitness values that vary exponentially with rank. This method, too, fails to attain the ESS.

8 Boltzmann Selection

In *Boltzmann selection*, a method inspired by the technique of *simulated annealing*, selection pressure is slowly increased over evolutionary time to gradually focus search [Mitchell, 1996]. Given a fitness of f, Boltzmann selection assigns a new fitness, f', according to the differentiable function:

$$f' = e^{\beta f} \tag{5}$$

where $\beta > 0$ and higher values of β give higher selection pressure.

Agents are then selected in proportion to their new fitnesses, f'. In contrast to the selection methods seen above, this selection method can maintain arbitrary fixed points without modification. But, it too exhibits multiple regimes of behavior. For low selection pressures, Boltzmann selection can preserve the ESS attractor. Figure 6 (top) shows the map of Boltzmann selection in the Hawk-Dove game with $\beta = 0.05$; we see that the ESS is intact. If we increase the selection pressure to $\beta = 0.2$, as in Figure 6 (middle), a true limit cycle results and the ESS becomes an unstable fixed point. A higher pressure of $\beta = 0.5$ brings the system to the edge of chaos, seen in Figure 6 (bottom), yielding a small but positive Liapunov exponent. The analogy between low annealing temperature and high selection pressure is strained—too low a "temperature" actually destabilizes the system.

9 Large Games

How much can we generalize from the results obtained with the two-strategy Hawk-Dove game? The discontinuities found in truncation, (μ, λ)-ES, and rank selection disrupt evolutionary dynamics regardless of the number of strategies because they are unable, in principle, to attain arbitrary fixed points. Specifically, these methods can neither attain nor maintain a Nash equilibrium that has more than one strategy in support, such as a polymorphic ESS.

For Boltzman selection, the destabilizing effect of high selection pressure is dampened as the number of strategies increases. The higher number of dimensions of large games tends to bring fitness values closer together. Thus, higher selection pressure is required to induce chaos in large games.

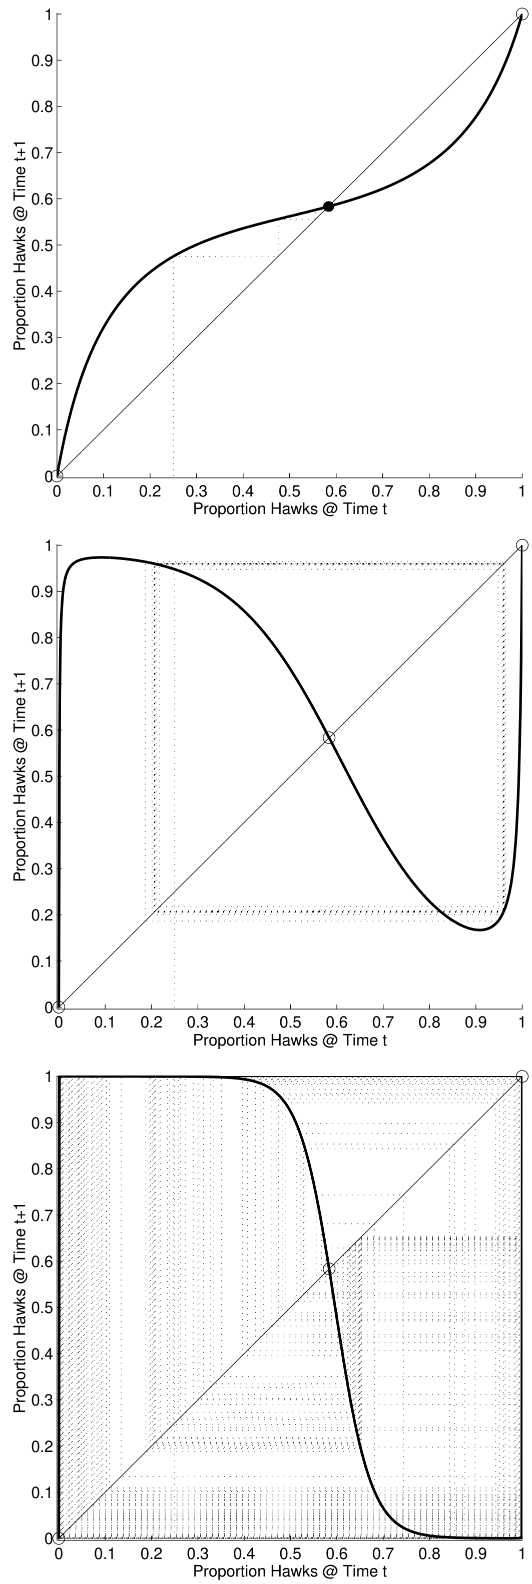

Figure 6: Map diagram of Boltzman selection method with $\beta = 0.05$ (top), 0.2 (middle), and 0.5 (bottom).

10 Implications for Coevolutionary Algorithms

Our results show that truncation, (μ,λ)-ES, and rank selection methods disrupt the "normal" dynamics of evolutionary game theory. Indeed, they introduce new dynamics that are incapable of having attractive polymorphic population states, *regardless of the payoff matrix*. The significance of these results is perhaps better appreciated once we recognize that evolutionary game theory corresponds to a generational *coevolutionary algorithm*, with fitness-proportionate selection, operating under "ideal" conditions: an infinite population, complete mixing (such that every agent plays against every other), noiseless payoffs, and complete knowledge of the phenotype space. Given the strategy set of any coevolutionary game, game theory will indicate the set of Nash equilibria (the "rational" strategies). *Evolutionary* game theory will highlight the subset of Nash equilibria that can be attained through fitness-proportionate differential reproduction (the evolutionary stable strategies). That truncation, (μ,λ)-ES, and rank selection all fail to converge onto polymorphic ESSs and fail to maintain polymorphic Nash equilibria under "ideal" conditions raises reasonable doubt that they will succeed under the less-than-ideal conditions of real world coevolutionary algorithms. Indeed, results obtained are likely to represent the peculiarities of the reproductive dynamical system more than any fundamental feature of the domain under study. We therefore conclude that these methods are probably inappropriate for use in single-population coevolutionary algorithms.

Boltzman selection, on the other hand, is capable of respecting the dynamics and equilibria of evolutionary game theory, provided the selection pressure is not too high. Indeed, any number of differentiable functions can be substituted for fitness proportional selection, and provide variable selection pressure while maintaining "proper" operation.

Empirical evidence in [Meuleau and Lattaud, 1995] and [Fogel and Fogel, 1995, Fogel et al., 1998] agrees with our findings with respect to (μ,λ)-ES and truncation selection, respectively. The former study notes dramatic differences between results obtained using (μ,λ)-ES and fitness-proportionate selection in coevolving strategies for the iterated prisoner's dilemma. The latter studies [Fogel and Fogel, 1995, Fogel et al., 1998] are testing the effects of finite populations, noisy payoffs, and incomplete mixing on the dynamics and equilibria of evolutionary game theory. They use the Hawk-Dove game as an experimental framework, but use truncation selection instead of the canonical fitness-proportional selection in their experiments. As we have seen, truncation selection is unable to reproduce the "correct" dynamics for the Hawk-Dove game, even under ideal conditions. The results of Fogel, et al, are consistent with ours, but they attribute the absence of normal Hawk-Dove dynamics to the factors under their study; because the ESS of $7/12^{ths}$ Hawks does not emerge, Fogel, et al, conclude that evolutionary game theory loses it predictive power under real-world conditions. We strongly suspect that the salient factor in their results is the use of truncation selection, and not finite populations, noisy payoffs, or incomplete mixing. Indeed, in light of our results, we believe that the many coevolutionary (and game-theoretic) investigations in the literature that use these selection methods may require a second look, especially where a single population is used and the domain under investigation is not a constant-sum game.

11 When does a map have a stable fixed-point?

We have analyzed the stability properties of four particular selection methods. In this section, we review a standard test that can be applied to any differentiable selection function to determine its stability properties [Easton, 1998]. 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.

Let $p_{t+1} = M(p_t)$ be a map with a fixed point at p^{fix}. We first linearize the system by calculating its first derivative at p^{fix}. For a game of n strategies, we have an $m = n-1$ dimensional map. If $n > 2$, then the map is a *multi-variable* function, and we need to calculate its *Jacobian Matrix*:

$$\partial M(p) = \begin{pmatrix} \frac{\partial M_{p^1}}{\partial p^1} & \cdots & \frac{\partial M_{p^1}}{\partial p^m} \\ \vdots & \ddots & \vdots \\ \frac{\partial M_{p^m}}{\partial p^1} & \cdots & \frac{\partial M_{p^m}}{\partial p^m} \end{pmatrix} \qquad (6)$$

where $\frac{\partial M_{p^i}}{\partial p^j}$ is the partial derivative of function variable p^i with respect to variable p^j.

The test for convergence is to check whether the eigenvalues of $\partial M(p^{\text{fix}})$ are within the interior of the unit-circle. That is, we check each eigenvalue, λ, of the Jacobian to see if $\|\lambda\| < 1$ (where λ is potentially a complex number). 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 one-dimensional map, such as the one we saw for fitness-proportionate selection in Figure 1. This map contains a stable fixed-point, such that if we iterate the map from a point near the fixed-point, we will converge onto the fixed-point. Why do points in the neighborhood of the fixed-point converge onto it? Take a point at an offset from the fixed point, $p_0 = p^{\text{fix}} + \mu$, close enough to the fixed point that we can treat the map as linear. The linearization of the map in the region allows us to approximate the map as $M(p) \approx p^{\text{fix}} + \lambda(p - p^{\text{fix}})$, 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, $p_1 = M(p_0) \approx p^{\text{fix}} + \lambda(p^{\text{fix}} + \mu - p^{\text{fix}}) = p^{\text{fix}} + \lambda\mu$. If $|\lambda|$ is less than 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, $p_n = p^{\text{fix}} + \lambda^n\mu$, and (if $|\lambda| < 1$) bring it closer and closer to the fixed-point.

Thus, the one-dimensional map test is whether the absolute value of the derivative (slope) at the fixed point is less than one. The multi-dimensional test is based on the same convergence properties as the one-dimensional case. In a way, taking the eigenvalues of the Jacobian matrix is equivalent to breaking the multi-dimensional system down into constituent one-dimensional systems, where each eigenvalue represents the rate of change (derivative) of each one-dimensional degree of freedom of the multi-dimensional 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 multi-dimensional case, the test becomes whether all the eigenvalues' magnitudes are less than one, i.e., in the unit circle.

12 Conclusion

We investigate the game-theoretic properties of selection methods commonly used in evolutionary algorithms to see if they are able to maintain the dynamics and equilibria of evolutionary game theory. Using the Hawk-Dove game as a backdrop, we substitute the fitness-proportionate selection method assumed in evolutionary game theory with these other selection schemes. We find that truncation, (μ, λ)-ES, and linear rank selection are unable to maintain polymorphic evolutionary stable strategies. Instead, these methods exhibit behaviors such as limit cycles, chaos, or fixed points that are essentially unrelated to the values of the payoff matrix. Boltzman selection, however, is able to retain the dynamics and equilibria seen in evolutionary game theory, provided that the selection pressure is not too high.

These results transcend the field of evolutionary game theory, however. Because of the correspondence between evolutionary game theory and coevolutionary algorithms, our results indicate that selection methods cannot be moved wholesale from evolutionary to *co*evolutionary frameworks without careful consideration. Specifically, three of the four methods we consider in this paper appear to be pathological in the context of single-population coevolution in non-zero sum games. Nevertheless, a test from dynamical systems theory allows one to determine analytically the appropriateness for coevolution of specific differentiable selection functions, without the need for empirical investigation.

13 Acknowledgments

The authors wish to thank the members of the DEMO Lab for their support of our research.

Bibliography

[Alliot et al., 1995] Alliot, J.-M. et al., editors (1995). *Aritificial Evolution (AE 95)*. Springer-Verlag.

[Bäck, 1995] Bäck, T. (1995). Evolution strategies: An alternative evolutionary algorithm. In [Alliot et al., 1995], pages 3–20.

[Dugatkin and Reeve, 1998] Dugatkin, L. A. and Reeve, H. K., editors (1998). *Game Theory and Animal Behavior*. Oxford University Press.

[Easton, 1998] Easton, R. W. (1998). *Geometric methods for discrete dynamical systems*. Oxford University Press.

[Fogel, 1997] Fogel, D. B. (1997). An overview of evolutionary programming. In Davis, L. D., De Jong, K., Vose, M. D., and Whitley, L. D., editors, *Evolutionary Algorithms*, pages 89–109. Springer.

[Fogel and Fogel, 1995] Fogel, D. B. and Fogel, G. B. (1995). Evolutionary stable strategies are not always stable under evolutionary dynamics. In *Evolutionary Programming IV*, pages 565–577.

[Fogel et al., 1998] Fogel, G. B., Andrews, P. C., and Fogel, D. B. (1998). On the instability of evolutionary stable strategies in small populations. *Ecological Modelling*, 109:283–294.

[Fudenberg and Levine, 1998] Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games*. MIT Press.

[Goldberg and Deb, 1991] Goldberg, D. E. and Deb, K. (1991). A comparitive analysis of selection schemes used in genetic algorithms. In Rawlins, G. J., editor, *Foundations of Genetic Algorithms (FOGA 1)*, pages 69–93.

[Hancock, 1994] Hancock, P. J. (1994). An empirical comparison of selection methods in evolutionary algorithms. In Fogarty, T. C., editor, *Evolutionary Computing (AISB '94)*, pages 80–94.

[Hofbauer and Sigmund, 1998] Hofbauer, J. and Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press.

[Maynard-Smith, 1982] Maynard-Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press.

[Meuleau and Lattaud, 1995] Meuleau, N. and Lattaud, C. (1995). The artificial evolution of cooperation. In [Alliot et al., 1995], pages 159–180.

[Mitchell, 1996] Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. MIT Press.

[Samuelson, 1997] Samuelson, L. (1997). *Evolutionary Games and Equilibrium Selection*. MIT Press.

[Weibull, 1995] Weibull, J. (1995). *Evolutionary Game Theory*. MIT Press.