

A Game-Theoretic Investigation of Selection Methods in Two-Population Coevolution

Sevan G. Ficici

Division of Engineering and Applied Sciences
Harvard University
Cambridge, Massachusetts 02138 USA
sevan@eecs.harvard.edu

ABSTRACT

We examine the dynamical and game-theoretic properties of several selection methods in the context of two-population coevolution. The methods we examine are fitness-proportional, linear rank, truncation, and (μ, λ) -ES selection. We use simple symmetric variable-sum games in an evolutionary game-theoretic framework. Our results indicate that linear rank, truncation, and (μ, λ) -ES selection are somewhat better-behaved in a two-population setting than in the one-population case analyzed by Ficici et al. [4]. These alternative selection methods maintain the Nash-equilibrium attractors found in proportional selection, but also add non-Nash attractors as well as regions of phase-space that lead to cyclic dynamics. Thus, these alternative selection methods do not properly implement the Nash-equilibrium solution concept.

Categories and Subject Descriptors

I.2.8 [Computing Methodologies]: Artificial Intelligence—*Problem Solving, Control Methods, and Search*; I.2.6 [Computing Methodologies]: Artificial Intelligence—*Learning*

General Terms

Theory

Keywords

Coevolution, selection method, population dynamics, solution concept

1. INTRODUCTION

Recent work by Ficici et al. [4] examines the game-theoretic and dynamical properties of several selection methods in the context of single-population coevolution. They use simple 2x2 symmetric two-player games that have polymorphic Nash equilibria. In particular, they focus on poly-

morphic Nash equilibria that are attractors of the canonical replicator dynamics [10], in which agents are selected to reproduce in proportion to fitness. Though proportional selection has been extensively studied from the perspective of evolutionary game theory [11, 10], the many alternative selection methods that are used in evolutionary algorithms have not. Ficici et al. [4] show that most alternatives to proportional selection are unable to converge onto polymorphic Nash equilibria; instead, they find a variety of other dynamics, including convergence to non-Nash fixed-points, cyclic dynamics, and chaos.

In this paper, we broaden the scope of their investigation by looking at 2x2 symmetric variable-sum games in the context of two-population coevolution. We consider both classes of 2x2 games that have polymorphic Nash equilibria. We examine the behaviors of linear rank, truncation, and (μ, λ) -ES selection, and contrast them to the dynamics of proportional selection. Our empirical results show that the alternative selection methods maintain the attractors found in proportional selection, but also add non-Nash attractors as well as regions of phase-space that lead to cyclic dynamics. Thus, these alternative selection methods do not properly implement the Nash-equilibrium solution concept.

The paper is organized as follows. Section 2 defines the classes of games that we will use in our investigation. Section 3 details how we embed these games into a two-population coevolutionary framework. Section 4 describes how populations represent Nash equilibria. Section 5 completes the description of the evolutionary framework. Sections 6 through 9 describe our results, and Section 10 provides concluding thoughts.

2. 2X2 SYMMETRIC GAMES

We examine 2x2 symmetric variable-sum games for two players. A two-player game is *symmetric* when both players have available the same set of strategies and the identity (i.e., Player 1 or Player 2) of the player playing a particular strategy has no effect on the payoffs obtained. Examples of such games include the Hawk-Dove game [11] and the Prisoner's Dilemma [1]. A generic payoff matrix for a two-strategy symmetric variable-sum game for two-players is given by Equation 1. By convention, the payoffs are for the row player; thus, if one player uses Strategy X and the other uses Strategy Y, then the X-strategist earns a payoff of b and the Y-strategist earns c . (Since we are dealing with symmetric games, the column and row players are interchangeable.) We can divide the space of possible 2x2 games into three mutually exclusive and exhaustive classes,

as follows. (We will find that the first two classes hold the most interest for us, below.)

$$\mathbf{G} = \begin{array}{c|cc} & \text{X} & \text{Y} \\ \hline \text{X} & a & b \\ \text{Y} & c & d \end{array} \quad (1)$$

2.1 Game Type A

The 2x2 games in this class have payoffs $a > c$ and $b < d$; these games are known as *coordination games* [7], since the players are always better off if they play the same pure strategy (i.e., coordinate). These coordinated configurations (both playing X or both playing Y) are Nash equilibria in pure strategies: neither player has incentive to deviate unilaterally from the pure strategy it is using. For example, if the column player uses strategy X, then the row player should also use X; if the column player instead uses Y, then so too should the row player. Coordination games also have a mixed-strategy Nash equilibrium where both players randomize over X and Y with a particular distribution (which we will discuss below).

2.2 Game Type B

The games in this class have payoffs $a < c$ and $b > d$; an example of this type of game is the Hawk-Dove game [11]. Games in this class have a single Nash equilibrium in pure strategies: one player uses the X strategy and the other uses Y; given this configuration, neither player has incentive to deviate unilaterally. There also exists a mixed-strategy Nash equilibrium where both players use the same mixed strategy (we discuss this below, as well).

2.3 Game Type C

The final class of 2x2 games has payoffs $a < c$ and $b < d$. Here, Strategy X is the preferred strategy for a player regardless of what the other player does. Thus, we say that X *dominates* Y. A weaker form of domination can be obtained if one of the inequalities is replaced by equality. (Note that the payoff structure $a > c$ and $b > d$ is structurally identical, except that now Strategy Y is preferred.)

3. TWO-POPULATION SYSTEMS

Following Ficici et al. [4], we will use infinitely large populations of agents; we depart from this earlier work by examining two-population systems. Each agent in each population will play one of the two pure strategies of the game (X or Y); agents are not allowed to play mixed strategies (probability distributions over the pure strategies). Because we have only two pure strategies, we can represent the state of our two-population system with two variables. Let p represent the proportion of X-strategists in Population 1, and q the proportion of X-strategists in Population 2.

We assume complete mixing, that is, every agent in Population 1 interacts with every agent in Population 2 and *vice versa*. Agents accumulate payoffs as they interact with each other. Let w_{X_1} and w_{Y_1} be the cumulative payoffs received by X- and Y-strategists, respectively, in Population 1; similarly, let w_{X_2} and w_{Y_2} be the cumulative payoffs for individuals in Population 2. Given the population states p and q , we can calculate cumulative payoffs with the linear equations in Equation 2.

$$\begin{aligned} w_{X_1} &= qa + (1 - q)b \\ w_{Y_1} &= qc + (1 - q)d \\ w_{X_2} &= pa + (1 - p)b \\ w_{Y_2} &= pc + (1 - p)d \end{aligned} \quad (2)$$

To gain some insight into the dynamics we might expect, we can plot how the cumulative payoffs vary over the two-dimensional phase space. Figure 1 provides such a plot for example games of types A (top) and B (bottom). The horizontal span of an arrow indicates the difference $w_{X_1} - w_{Y_1}$, while the vertical span indicates the difference $w_{X_2} - w_{Y_2}$; thus each arrow indicates, for a particular location in phase space, the strategy that is most advantageous (and by how much) with respect to each population. The pure-strategy Nash equilibria are indicated by the solid circles, and the mixed-strategy Nash equilibrium is indicated by the open circle. This mixed-strategy Nash equilibrium is also known as a *polymorphic Nash equilibrium*.

4. MIXED STRATEGIES AND POLYMORPHIC POPULATIONS

Though individual agents cannot play mixed strategies, we can view a population as a whole to play one. We interpret the proportion p of X-strategists in Population 1 as specifying a mixed strategy where the probability of playing X is p and Y is $(1 - p)$; similarly, the proportion q of X-strategists in Population 2 specifies a mixed strategy where the probability of playing X is q and Y is $(1 - q)$. Given a two-player game, we may view Populations 1 and 2 as representing mixed strategies for Players 1 and 2, respectively.

The *support* of a mixed strategy is the set of pure strategies played with probability greater than zero. (Note that we may view a pure strategy as a degenerate mixture.) A population that contains more than one pure strategy is termed *polymorphic*; otherwise, the population is *monomorphic*.

In a symmetric game with a Nash-equilibrium mixed-strategy m , if Player 1 plays m , then the highest payoff obtainable by Player 2 is received by also playing m . But Player 2 has other choices, as well. In particular, if Player 2 plays any pure strategy in support of m , the same payoff is received [8]. Thus, the strategies in support of a Nash-equilibrium mixed-strategy will be at fitness equilibrium [11]. For game types A and B, above, this means that Population 1 is playing the Nash mixture when $w_{X_2} = w_{Y_2}$ (a similar statement applies to Population 2); we need merely solve for p to discover the proportion of X-strategists needed to achieve fitness equilibrium (a similar statement applies to q), as shown in Equation 3. Since the game is symmetric, we have $p = q$ when both populations are playing the Nash mixture.

$$p_{\text{eq}} = \frac{d - b}{a - c + d - b} \quad (3)$$

5. REPLICATION

Once cumulative payoffs are obtained, we create the next generation of the population. Following conventional evolutionary game theory [11], offspring are generated asexually and are clones of their parents; that is, we do not use variation operators. As a result, strategies absent from the

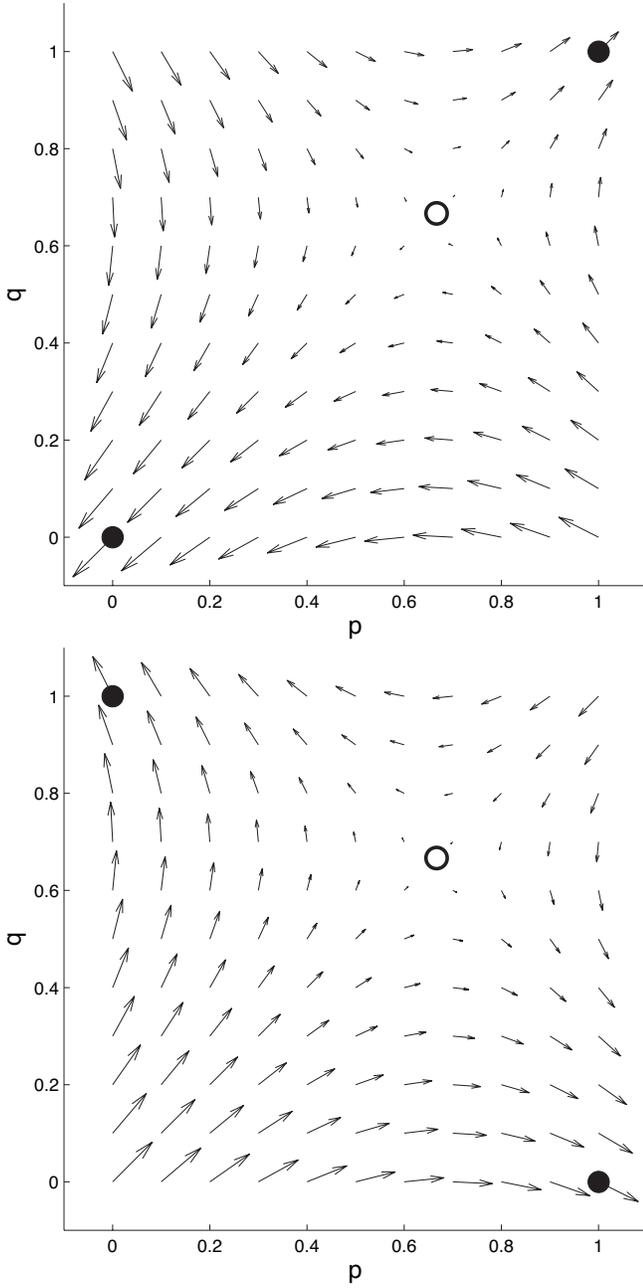


Figure 1: Variation of cumulative payoffs in two-dimensional phase space for Type-A (top) and Type-B (bottom) games. The horizontal span of an arrow indicates the difference $w_{X_1} - w_{Y_1}$; the vertical span indicates the difference $w_{X_2} - w_{Y_2}$. For example, in the Type-B game, the arrow that originates at $p = 1, q = 0$ indicates that when Strategies X and Y each play a population of all Y-Strategists ($q = 0$), then Strategy X strongly out-scores Strategy Y; further, when Strategies X and Y each play a population of all X-Strategists ($p = 1$), then Strategy Y weakly out-scores Strategy X. The arrows do not locate the point in phase space to which the two populations will go; the arrows merely indicate which pure strategy is most advantageous (and by how much) with respect to each population's state.

initial population cannot appear later in time. The way we calculate the number of offspring each agent receives is determined by the selection method we use. We next discuss the dynamical and game-theoretic properties of various selection methods, beginning with a conventional replication dynamic that is used in evolutionary game theory.

6. PROPORTIONAL SELECTION

Here we review the properties of proportional selection when applied to symmetric games in a two-population framework [10]. Given states p and q for Populations 1 and 2, respectively, Equation 4 shows how to calculate the next population states $f(p)$ and $g(q)$; individuals generate offspring in proportion to their fitness (i.e., cumulative payoff).

$$\begin{aligned} f(p) &= \frac{w_{X_1} p}{w_{X_1} p + w_{Y_1} (1 - p)} \\ g(q) &= \frac{w_{X_2} q}{w_{X_2} q + w_{Y_2} (1 - q)} \end{aligned} \quad (4)$$

6.1 Type-A Games

Figure 2 (top) illustrates the phase plot of proportional selection on a Type-A game, defined above. The payoffs we use in this example are $a = 3, b = 2, c = 2, d = 4$ (we use this game as our Type-A exemplar throughout the paper); using Equation 3, we find the Nash-equilibrium mixed-strategy to be at $p = q = 2/3$ (open circle). The dynamical system has two point-attractors, one at $p = q = 0.0$ and the other at $p = q = 1.0$ (closed circles); these attractors correspond to the “coordinated” pure-strategy Nash equilibria. The basin of attraction for the latter attractor is smaller as a result of the asymmetry in the payoffs. The mixed-strategy Nash equilibrium is an unstable fixed-point. Thus, if we begin the system at $p = q = 2/3$, then it will remain there; if we perturb either one of the population states (or both), then the system will diverge from the unstable fixed-point and move to one of the two attractors.

We can collapse this two-dimensional system down to one dimension by setting the initial condition such that $p = q$; this gives us the dynamics of a single population. Since all three Nash equilibria lie on the diagonal $p = q$, we know that the dynamics of the Type-A game remain essentially unchanged as we move from one population to two.

6.2 Type-B Games

Figure 2 (bottom) shows the phase plot of proportional selection on a Type-B game. The payoffs here are $a = 2, b = 4, c = 3, d = 2$ (we use this game as our Type-B exemplar throughout the paper); this is identical to the Type-A game, above, except we have switched the rows of the payoff matrix. Note that the mixed Nash equilibrium remains an unstable fixed-point at $p = q = 2/3$. The two pure-strategy Nash equilibria, however, have moved off of the $p = q$ diagonal; now, these point-attractors are found at $p = 0, q = 1$ and $p = 1, q = 0$. Though we still have two basins of attraction, note that the basins now have equal size; all points above the $p = q$ diagonal go to $p = 0, q = 1$, points below go to $p = 1, q = 0$. If we begin the system on the diagonal $p = q$ (collapsing it to the single-population case), we converge to the mixed-strategy equilibrium (i.e., both populations are polymorphic and have $2/3$ X-strategists). The mixed Nash equilibrium is a saddle-point of the two-population system, and the $p = q$ line is the stable (attracting) manifold.

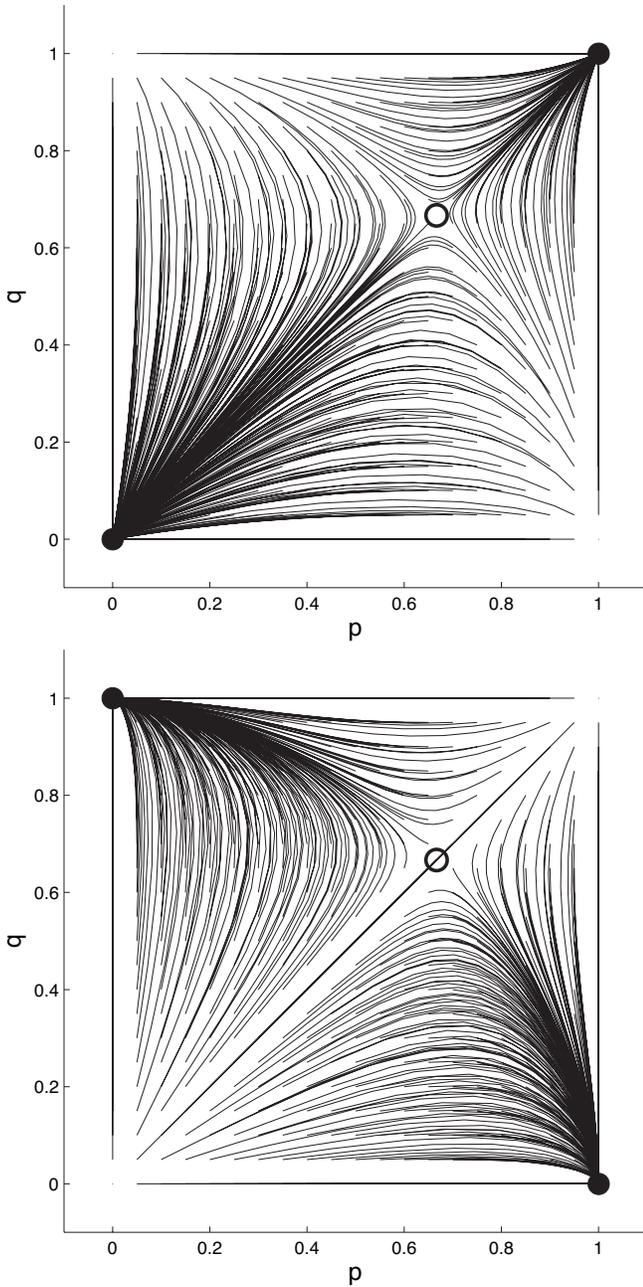


Figure 2: Proportional selection in Type-A games (top) and Type-B games (bottom). Each line traces a trajectory in phase space, with most trajectories leading to one of the pure-strategy Nash equilibria (solid circles). Note that each game-type includes two non-Nash fixed-points where both populations are monomorphic. For example, in the Type-A game at $p = 0, q = 1$, we know from Figure 1 that X-Strategists out-score Y-Strategists when they play against Population 2 at $q = 1$. Nevertheless, Population 1 does not contain any X-strategists at $p = 0$, and so its state cannot change; a similar statement applies to the fact that Y-Strategists out-score X-Strategists when they play against Population 1 at $p = 0$. The two populations lack a common strategy on which they can coordinate.

For Type-B games, therefore, we find that the dynamics do change as we move from one population to two. This result holds relevance to practitioners; the mere addition of a second population changes the solutions the system can deliver. Given a single population, the mixed-strategy Nash equilibrium (polymorphic population) is the single attractor; this outcome represents the Nash equilibrium where Players 1 and 2 adopt the same mixed strategy. When two populations are involved, the Nash-equilibrium mixed strategy is unstable, and one of two pure-strategy Nash equilibria will emerge instead; one population will converge to all-X while the other converges to all-Y. Note that a single population cannot represent such outcomes. For example, the Nash equilibrium where Player 1 adopts Strategy X and Player 2 adopts Strategy Y specifies that each of the two players uses a different pure strategy. A single population cannot simultaneously *represent* two distinct (pure or mixed) strategies, one for each player of the game. A population may *contain* two distinct pure strategies, but such a polymorphic population represents a single mixed strategy.

7. LINEAR RANKING

Linear ranking is an alternative to proportional selection that is often used in evolutionary computation [9, 12]. The general process begins by sorting the individuals in a population according to their cumulative scores; then, the individual(s) with the lowest cumulative score receives a rank value of 1.0, the individual(s) with the next highest cumulative score receives a rank value of 2.0, and so on. Individuals then reproduce in proportion to their rank values. For our simple 2x2 games, each agent playing the lower-scoring strategy in the population receives a rank value of 1.0, while agents playing the other strategy receive a rank value of 2.0.

7.1 Type-A Games

With Type-A games, linear ranking in a single-population will always converge onto one of the two pure-strategy Nash equilibria [4]. With the addition of a second population, however, we find that period-two cycles are also possible. Figure 3 (top) shows a phase plot for linear ranking in our example Type-A game. We notice that the trajectories are less smooth. The filled squares indicate initial conditions that lead to period-two cycles. The cycle-inducing region surrounds the stable (attracting) manifold of the saddle point located at the mixed Nash equilibrium (open circle). A cyclic trajectory oscillates across the $p = q$ diagonal.

7.2 Type-B Games

With Type-B games, linear ranking in a single population will always produce cyclic dynamics [4]. With the addition of a second population, we find that linear ranking is somewhat better behaved. Figure 3 (bottom) shows a phase plot for linear ranking in our example Type-B game. We notice that most initial conditions manage to converge to one of the two pure-strategy Nash equilibria. On the $p = q$ diagonal, however, we know that the single-population results will emerge; period-two cycles will be found. We find that the cycle-inducing region (filled squares) is not only on the diagonal, but also surrounds it. In Type-B games, the mixed Nash equilibrium is again a saddle point, but now the $p = q$ diagonal is the stable (attracting) manifold. Now, a cyclic trajectory oscillates across the unstable manifold (which is stable in Type-A games).

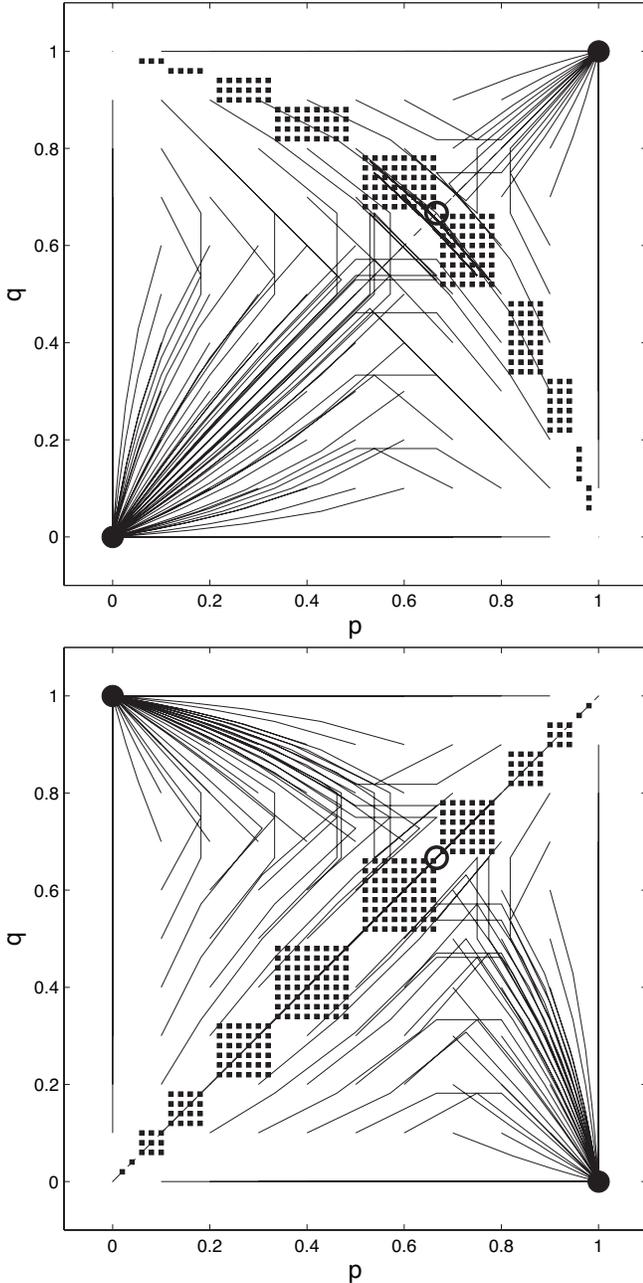


Figure 3: Linear ranking selection in Type-A games (top) and Type-B games (bottom). We sample the phase space at regular intervals to obtain our initial conditions; filled squares indicate initial conditions that lead to cyclic dynamics. For example, the initial condition $p = 0.60, q = 0.51$ has a cyclic orbit in the Type-B game. At this initial condition, Strategy X is the higher-scoring strategy in both populations; thus, the fitness of X and Y is 2 and 1, respectively. The next population state is $p = 0.75, q \approx 0.6755$ (this is across the unstable manifold from $p = 0.60, q = 0.51$); now, Strategy Y is the higher-scoring strategy in both populations. At this point, linear rank selection reverses the fitness values and so brings us back to the initial condition.

8. TRUNCATION

Truncation selection is often used in evolutionary programming [5] and the particular form we examine can be found in [6], for example. In this selection process, we begin by sorting the individuals in a population according to their cumulative scores. Each individual in the highest-scoring γ fraction of the population will create one offspring to replace one individual in the lowest-scoring γ fraction of the population; thus, γ expresses selection pressure. For example, let us suppose that $w_{X_1} > w_{Y_1}$, $p = 0.6$, and $\gamma = 0.1$. The highest-scoring tenth of the population is entirely X-strategists and the worst-scoring tenth is all Y-strategists; truncation will cause p to increase from 0.6 to 0.7.

8.1 Type-A Games

For Type-A games, one of the two pure-strategy Nash equilibria will always be obtained under truncation selection, given a single-population system [4]. We find that, with a two-population system, this is not always the case. Indeed, the behavior of truncation selection can be quite similar to that of linear ranking. Figure 4 (top) shows the phase diagram for our example Type-A game when $\gamma = 0.15$. We see that most initial conditions converge to one of the two pure-strategy Nash equilibria; there also exists a region of initial conditions (filled squares), surrounding the stable manifold, that lead to cyclic dynamics.

8.2 Type-B Games

For Type-B games, truncation selection cannot converge onto a Nash equilibrium, given a single population, regardless of $\gamma > 0$ [4]. With the addition of a second population, however, the two pure-strategy Nash equilibria become accessible from most initial conditions. Like linear rank selection, there exists a region of initial conditions surrounding the stable manifold that lead to cyclic dynamics. Figure 4 (bottom) shows the phase plot for our example Type-B game, when $\gamma = 0.15$.

If we set $\gamma = 0.5$, we know that virtually all initial conditions on the $p = q$ diagonal will cause the system to converge to $p = 1.0, q = 1.0$, which is not a Nash equilibrium [4]. We find that there exists, in addition, a region of initial conditions (filled squares) surrounding the stable manifold ($p = q$ diagonal) that also converge to this non-Nash fixed-point. The convergence to $p = 1.0, q = 1.0$ is easily explained. For example, let both p and q be between $1/2$ and $2/3$; with these population states, it is the case that X-strategists have higher cumulative payoffs than Y-strategists. Since X-strategists form at least half of the population, the next generation will necessarily be composed entirely of X-strategists. Thus, the largest box on the diagonal represents the critical region in which we transition to $p = 1.0, q = 1.0$ in the very next time-step. The remaining initial conditions belong to the iterated *pre-image* of this critical region; that is, they are initial conditions that will eventually fall into the critical region.

9. (μ, λ) -ES

(μ, λ) -ES selection is often used in evolution strategies [2]. Given our simple framework, which lacks variation mechanisms, (μ, λ) -ES reduces to a variation of the truncation method we examine above. As with truncation, the selection-pressure parameter γ indicates the fraction of the highest-scoring individuals in a population that will create offspring;

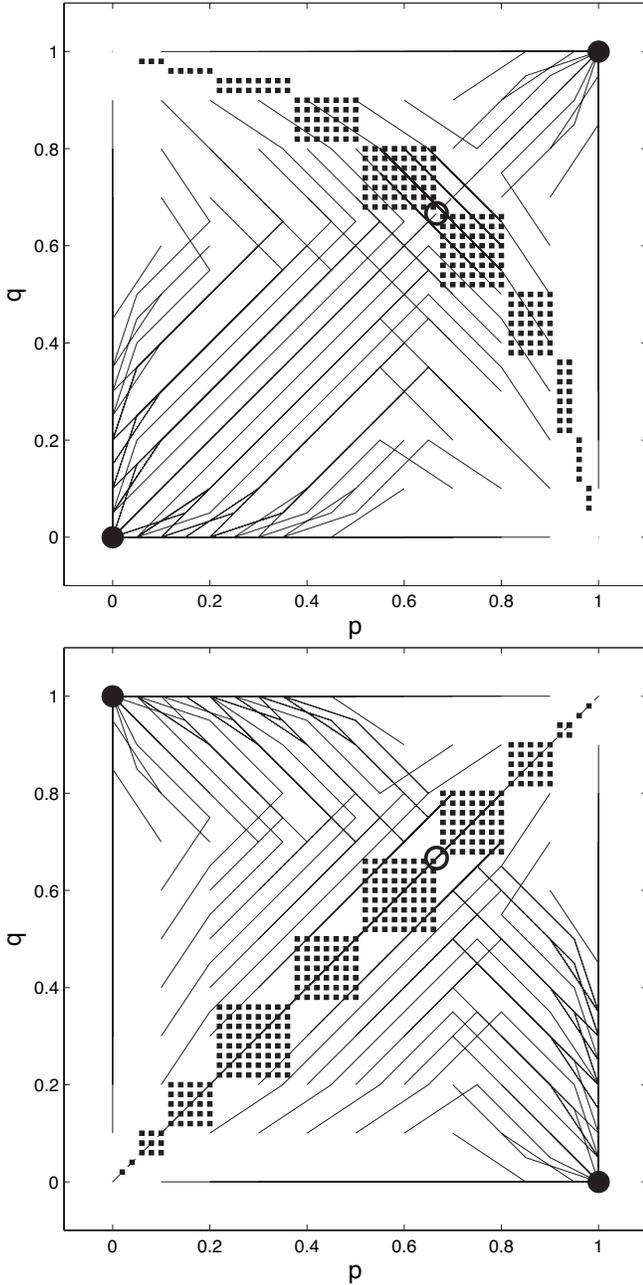


Figure 4: Truncation selection in Type-A games (top) and Type-B games (bottom), with $\gamma = 0.15$. We sample the phase space at regular intervals to obtain our initial conditions; filled squares indicate initial conditions that lead to cyclic dynamics. As with linear rank selection, the cycles obtained under truncation selection involve discontinuities.

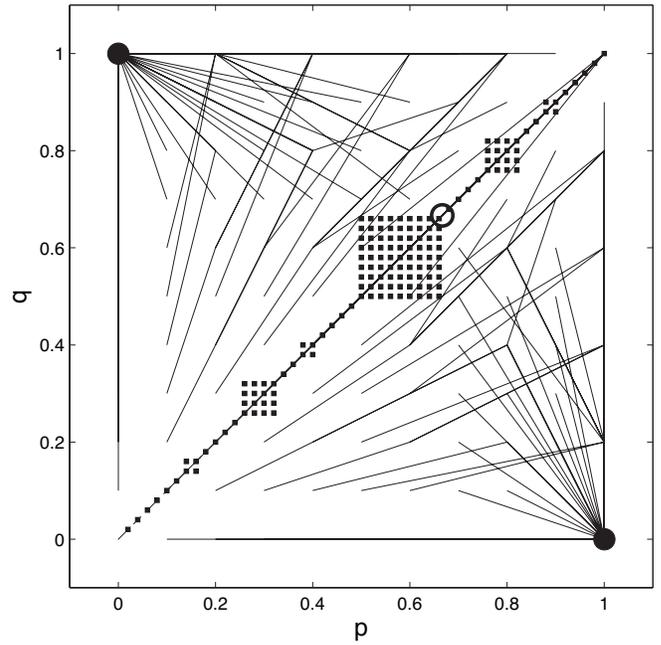


Figure 5: Truncation selection in Type-B game with $\gamma = 0.5$. We sample the phase space at regular intervals to obtain our initial conditions; filled squares indicate initial conditions that lead to a non-Nash attractor at $p = 1.0, q = 1.0$.

rather than just replace the lowest-scoring γ fraction of the population, we create enough offspring to replace the entire population. For example, let us suppose that $w_{X_1} > w_{Y_1}$, $p = 0.1$, and $\gamma = 0.2$. The highest-scoring two-tenths of the population is composed of all the X-strategists (the top tenth) and an equal proportion of Y-strategists (the next-best tenth). The subsequent generation will therefore be composed of equal numbers of X- and Y-strategists (i.e., p will move from 0.1 to 0.5).

9.1 Type-A Games

With Type-A games, one of the two pure-strategy Nash equilibria will be obtained, given a single population, regardless of the selection pressure $\gamma > 0$ [4]. We find that, with a two-population system, (μ, λ) -ES selection can also find these Nash equilibria, provided that selection pressure is not too severe. For high selection pressures (i.e., low γ), (μ, λ) -ES creates non-Nash attractors that we have not seen above. While most initial conditions do converge to a Nash equilibrium, many do not. Figure 6 (top) shows the phase plot for our example Type-A game with $\gamma = 0.1$; the figure indicates initial conditions (filled squares) that converge to the fixed-point $p = 0.0, q = 1.0$. This fixed-point represents a non-coordinated configuration between the game players, where one player uses Strategy X while the other uses Y. Similarly, Figure 6 (bottom) shows the initial conditions that converge onto the fixed-point $p = 1.0, q = 0.0$, again a non-coordinated state.

9.2 Type-B Games

With Type-B games, (μ, λ) -ES cannot converge onto a Nash-equilibrium, given a single population, regardless of

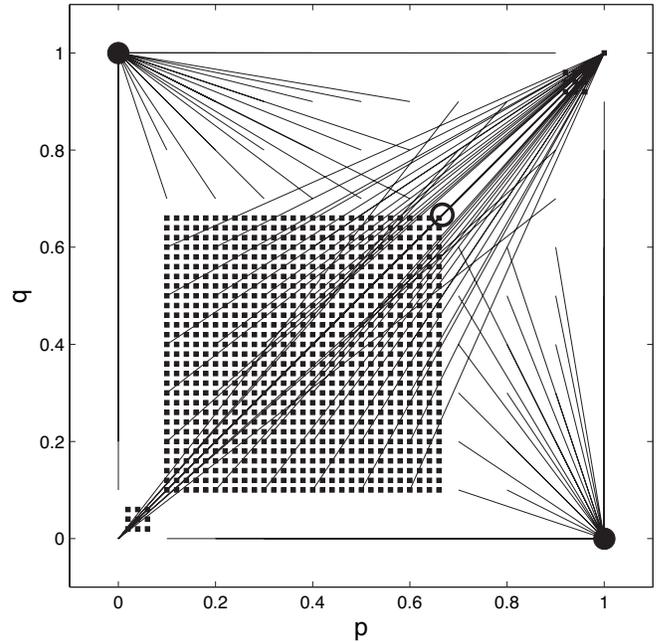
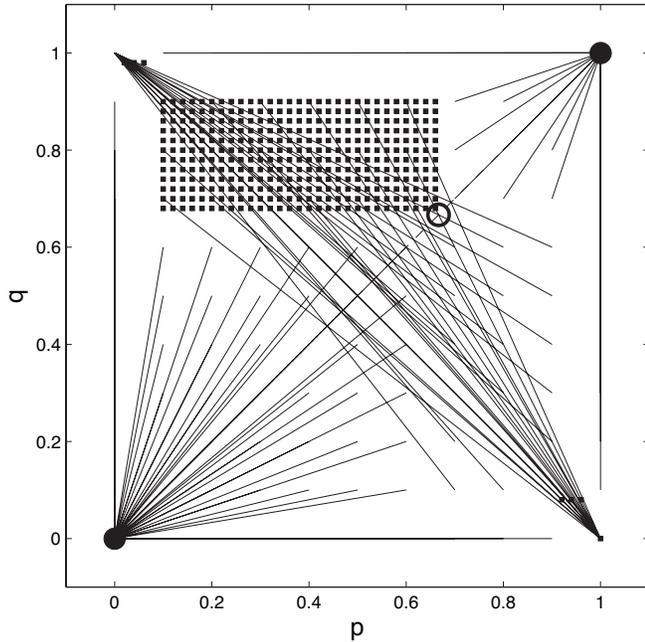
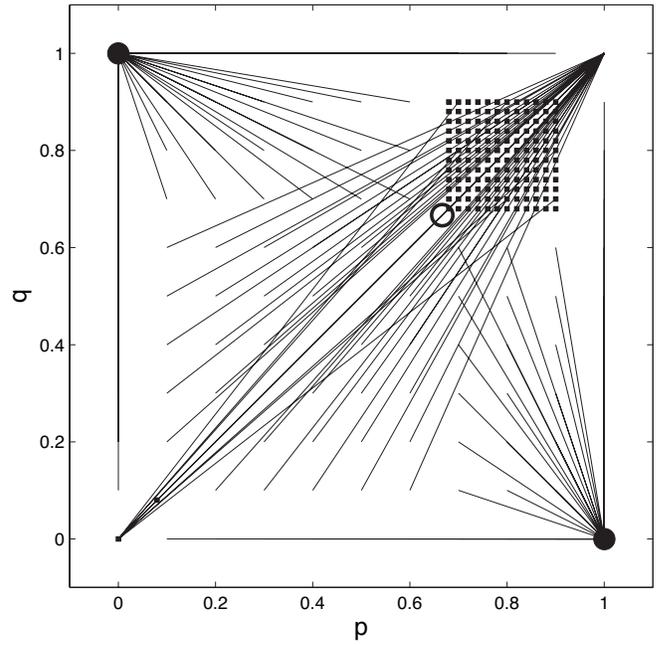
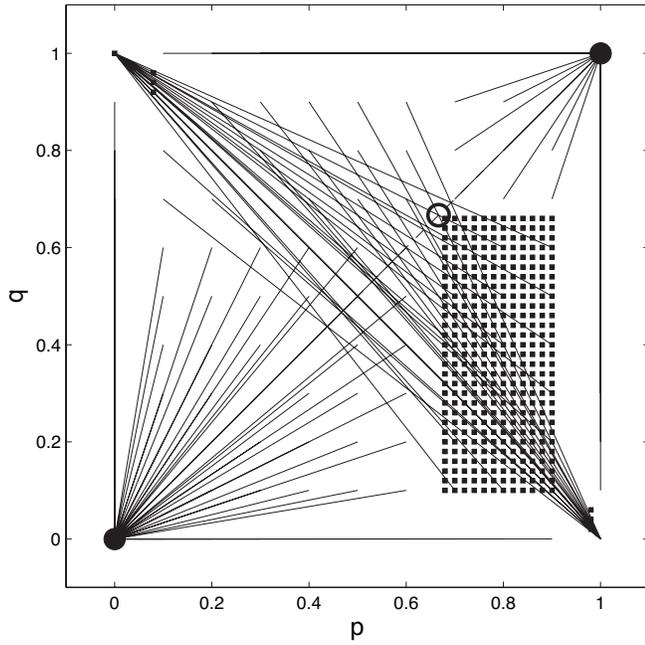


Figure 6: Type-A games under (μ, λ) -ES selection with $\gamma = 0.1$. We sample the phase space at regular intervals to obtain our initial conditions; filled squares indicate initial conditions that converge onto a non-Nash attractor. Top: Initial conditions that converge onto $p = 0.0, q = 1.0$. Bottom: Initial conditions that converge onto $p = 1.0, q = 0.0$.

Figure 7: Type-B games under (μ, λ) -ES selection with $\gamma = 0.1$. We sample the phase space at regular intervals to obtain our initial conditions; filled squares indicate initial conditions that converge onto a non-Nash attractor. Top: Initial conditions that converge onto $p = q = 0.0$. Bottom: Initial conditions that converge onto $p = q = 1.0$.

$\gamma > 0$ [4]. We find that with two populations, (μ, λ) -ES is capable of converging onto either of the pure-strategy Nash equilibria, again provided that selection pressure is not too severe. Figure 7 shows the phase diagram for our example Type-B game with $\gamma = 0.1$. While the pure-strategy Nash equilibria are still accessible from some initial conditions, we find that most of the phase space converges onto a non-Nash attractor. The top and bottom graphs in Figure 7 indicate the initial conditions (filled squares) that converge onto $p = q = 0.0$ and $p = q = 1.0$, respectively. Note that the initial conditions are on the opposite side of the saddle-point from the fixed-point they converge onto.

10. REVIEW AND CONCLUSION

We use a simple evolutionary game-theoretic framework to examine the dynamics of several selection methods—fitness-proportional, linear rank, truncation, and (μ, λ) -ES—in the context of two-population coevolution. We use simple (2x2) symmetric variable-sum games for two players. We focus on two classes of game. Games in either class have three Nash-equilibrium fixed-points in the two-dimensional phase-space; two of the Nash equilibria involve pure strategies while the third is a mixed-strategy Nash equilibrium.

The selection methods we consider have been analyzed by Ficici et al. [4] in the context of single-population coevolution; their results show that all selection methods will converge onto Nash equilibrium if the game belongs to the game class Type-A (see Section 2.1). In contrast, only fitness-proportional selection converges to Nash equilibrium if the game belongs to class Type-B (see Section 2.2); the other methods (that we consider in this paper) cannot converge onto Nash equilibrium. Instead a variety of other dynamics are reported; these include cyclic dynamics, non-Nash attractors, and chaos.

Our findings in this paper show that linear rank, truncation, and (μ, λ) -ES selection are better-behaved in a two-population setting than in the one-population case analyzed in [4]. The mixed-strategy Nash-equilibrium (in the interior of the phase space) is well known to be unstable with fitness-proportional selection [10]; the alternative selection methods we examine are no different in this respect. The pure-strategy Nash equilibria are known to be attractors of fitness-proportional selection [10], and they remain attractors for the alternative methods we examine.

Nevertheless, the alternative selection methods we consider do introduce some additional behaviors that are not found with fitness-proportional selection. Many initial conditions lead to cyclic dynamics, while others converge onto non-Nash fixed-points. Thus, one cannot use these selection methods with confidence if the desired solution concept is Nash equilibrium. On a more general note, our results further demonstrate that the mechanisms at work in a co-evolutionary algorithm (here, we look at selection methods) can profoundly affect the *de facto* solution concept that is implemented, and may cause it to diverge from the solution concept that we intend to implement [3].

While we use only two exemplar games in our investigation (one each from Type-A and Type-B), our results generalize across the games in these classes; results for different games will differ in numerical detail, but be qualitatively identical. Future work will expand our results by moving to higher dimensions (where the games have more than two strategies) and by moving to asymmetric (or bima-

trix) games (where the two populations use different strategy sets).

11. ACKNOWLEDGMENTS

The author thanks Anthony Bucci, Edwin de Jong, Jordan Pollack, Paul Wiegand, Shivakumar Viswanathan, and members of the DEMO Lab. Preparation of this paper was partially supported by DARPA’s Architectures for Cognitive Information Processing program grant NBCH-1-05-0005.

12. REFERENCES

- [1] R. Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.
- [2] T. Bäck. Evolution strategies: An alternative evolutionary algorithm. In J.-M. Alliot, E. Lutton, E. Ronald, M. Schoenauer, and S. D., editors, *Artificial Evolution (AE 95)*, pages 3–20. Springer-Verlag, 1995.
- [3] S. G. Ficici. *Solution Concepts in Coevolutionary Algorithms*. PhD thesis, Brandeis University, May 2004.
- [4] S. G. Ficici, O. Melnik, and J. B. Pollack. A game-theoretic and dynamical-systems analysis of selection methods in coevolution. *IEEE Transactions on Evolutionary Computation*, 9(6):580–602, 2005.
- [5] D. B. Fogel. An overview of evolutionary programming. In L. D. Davis, K. De Jong, M. D. Vose, and L. D. Whitley, editors, *Evolutionary Algorithms*, pages 89–109. Springer, 1997.
- [6] D. B. Fogel and G. B. Fogel. Evolutionary stable strategies are not always stable under evolutionary dynamics. In J. R. McDonnell, R. G. Reynolds, and D. B. Fogel, editors, *Evolutionary Programming IV: The Proceedings of Fourth Annual Conference on Evolutionary Programming*, pages 565–577. MIT Press, 1995.
- [7] D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*. MIT Press, 1998.
- [8] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1998.
- [9] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, 1989.
- [10] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.
- [11] J. Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982.
- [12] M. Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, 1996.