
Effects of Finite Populations on Evolutionary Stable Strategies

Sevan G. Ficici and Jordan B. Pollack
DEMO Lab—Computer Science Department
Volen National Center for Complex Systems
Brandeis University
Waltham Massachusetts USA
www.demo.cs.brandeis.edu

Abstract

A strong assumption made in evolutionary game theory (EGT) [7] is that the evolving population is infinitely large. Recent simulations by Fogel, et al, [3, 4, 5] show that finite populations produce behavior that, at best, deviate with statistical significance from the *evolutionary stable strategy* (ESS) predicted by EGT. They conclude that evolutionary game theory loses its predictive power with finite populations. In this paper, we revisit the question of how finite populations affect EGT dynamics. By paying particular attention to the operation of the selection mechanisms used by Fogel, et al, we are able to account for the divergence between ESS predictions (based on infinite populations) and results observed in our own finite-population simulations. We then show that Baker’s SUS [1] selection method corrects the divergence to a great extent. We thus conclude that the dynamics of EGT, and particularly ESSs, can indeed apply to finite-population systems.

1 Introduction

The primary contribution of evolutionary game theory (EGT) is the concept of the *evolutionary stable strategy* (ESS) [7]. The ESS is a refinement of Nash equilibrium that does not require agents to be rational to attain it. Rather, agents achieve equilibrium through a process of Darwinian selection. At least three strong assumptions are made in the EGT formalism. First, the evolving population is assumed to be infinitely large. Second, the payoffs that agents receive are assumed to be without noise. Third, each agent is assumed to play against every other agent to determine its fitness—there is *complete mixing*. Recently, Fogel,

et al, have questioned the usefulness of evolutionary game theory in real-world situations where these assumptions, particularly the first, do not hold [3, 4, 5]. They begin their simulations with the population precisely at the ESS, and discover that the population consistently moves away from it. At best, their simulation results differ from theoretical ESS values with statistical significance; at worst, their results bear no semblance to the ESS whatsoever. Thus, they conclude that evolutionary game theory loses predictive power once these assumptions are relaxed.

In this paper, we concentrate on the first (and perhaps strongest) of the above assumptions and revisit the question of whether ESS dynamics can exist in finite populations. Fogel, et al, report using two different selection methods, truncation, and proportional roulette-wheel selection. By paying particular attention to the operation of the selection mechanism, we are able to account for the divergence between ESS predictions (based on infinite populations) and results observed in our own finite-population simulations. We then examine Baker’s SUS selection method [1] that corrects the divergence to a great extent. We thus conclude that the dynamics of evolutionary game theory, and particularly the ESS, can indeed apply to finite-population systems. Further, the selection method used in a simulation can distort, or even disrupt completely, the dynamics we may expect to see.

We begin with a brief introduction to evolutionary game theory in Section 2, and then review and analyze relevant previous work in Section 3. In Section 4 we introduce the methodology used in our own finite-population simulations and give results. Section 5 analyzes the operation of our selection method and Section 6 constructs a method to predict the amount of divergence to expect for a particular population size. Section 7 examines the performance of Baker’s SUS selection method. Section 8 offers final remarks.

2 Evolutionary Game Theory

The Hawk-Dove game forms the backdrop for our investigation (as in [3, 4, 5]); it has two *pure strategies*, H (hawk) and D (dove). The *payoff matrix*, G , for this game is shown in Equation 1. Each entry, $E(i, j)$, in row i , column j , is the *expected value* of the payoff given to an agent playing strategy i when matched against an agent playing strategy j —payoffs are assumed to be without noise. The evolving population of agents is assumed to be infinitely large. The proportions with which the strategies of G are used in the population can be represented by a column vector, p ; the elements of p sum to 1.0. The *fitnesses*, f , of the strategies are determined by a weighted sum of the payoffs in G according to the proportions in p , and can be computed by matrix multiplication, as in Equation 2. This equation assumes *complete mixing*, that is, all agents play against all others. The next generation of agents is formed through *fitness-proportionate selection*—each strategy increases its representation, or “reproduces,” in direct proportion to its fitness. The reproductive process is described by the difference equation shown in Equation 3.

$$G = \begin{array}{c|cc} & \text{H} & \text{D} \\ \hline \text{H} & -25 & 50 \\ \text{D} & 0 & 15 \end{array} \quad (1)$$

$$f = G * p + w_0 \quad (2)$$

$$p' = \frac{p \times f}{p \bullet f} \quad (3)$$

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.

At the heart of evolutionary game theory is the concept of the evolutionary stable strategy. A population of agents that play an ESS cannot be “invaded” by a small number of agents that play some mutant strategy; hence, the population is evolutionarily stable. A *polymorphic* ESS actually involves two or more pure strategies, all of which receive the same fitness. The Hawk-Dove game has a polymorphic ESS where $7/12^{th}$ s of the population plays hawk ($p_H = \frac{7}{12}$), and $5/12^{th}$ s dove ($p_D = \frac{5}{12}$). According to Maynard-Smith, for a strategy (or polymorphic population), i , to be an ESS it must fulfill one of these two requirements against a mutant strategy, j , for all $j \neq i$ [7]:

$$E(i, i) > E(j, i) \quad \mathbf{OR} \quad (4)$$

$$E(i, i) = E(j, i) \quad \mathbf{AND} \quad E(i, j) > E(j, j) \quad (5)$$

3 Review of Previous Work

In this section, we review the experiments reported in [3, 4, 5], where Fogel, et al, examine various factors that may influence EGT dynamics. All experiments have finite populations and generational reproduction. Two different selection methods are used. The *truncation* selection method operates by sorting agents according to fitness and then replacing the worst v percent of the agents with copies of the best v percent. The value of v lies in the interval $[0, 50]$; higher values exert a higher selection pressure. The *fitness-proportionate roulette-wheel* selection method biases a random variable (the roulette wheel) in proportion to the agents’ fitnesses. The wheel is then “spun” n times to create an offspring population of size n .

3.1 Truncation Selection

In [3], truncation selection is used with maximum selection pressure, $v = 50\%$. This means that the worst half of the population is replaced by copies of the best half to form the next generation. Three experiments are described. Population sizes of 60, 600, and 6000 are used in each experiment and all pair-wise encounters occur during fitness evaluation (complete agent mixing). The first experiment adds a form of noise to the payoffs; the second introduces an additional, but slight, mutation bias (such that a hawk might become a dove, and vice versa). The third experiment is presented as a control—payoffs without noise and no mutation. In every case, the population is begun at the ESS and observed to move away. The results of the first two experiments are reported as essentially indistinguishable: for population sizes of 60 and 600, both give apparently chaotic behavior or a series of short-lived, quasi limit cycles. A population of size 6000, however, converges to all hawks in the first experiment (the second experiment always introduces some number of mutations). In the control experiment, the population is reported to jump immediately to all hawks. Given the clear absence of an ESS, Fogel, et al, conclude that the dynamics of evolutionary game theory not only assume, but *require* an infinite population.

Nevertheless, EGT also assumes fitness-proportional selection—not truncation. Let us consider the behavior of truncation selection in the case of noiseless payoffs. At the ESS, both strategies (and hence all agents) receive the same fitness: $p_H = \frac{7}{12} \Rightarrow f_H = f_D$. In contrast, $p_H < \frac{7}{12} \Rightarrow f_H > f_D$, and $p_H > \frac{7}{12} \Rightarrow f_H < f_D$. This describes a simple feedback mechanism (that regulates properly with proportional selection). Let us consider a population where the proportion of hawks is $.5 \leq p_H < \frac{7}{12}$. Because the proportion of hawks is

below the ESS, the hawks will receive higher fitness than the doves. Yet, the hawks comprise at least 50% of the population. Thus, the top 50% of the population can only contain hawks, and so, with truncation selection, the next generation must contain 100% hawks. This is true regardless of the size of the population—indeed, *even an infinite population*. At the ESS itself, the behavior of truncation selection is ill-defined because all agents have the same fitness—who are in the top 50%? Unless special care is taken to account for ties, the fixed point of all hawks will again emerge. The results of the third experiment (control) are consistent with this analysis.

Now let us consider the first experiment, which adds the following noise to the payoffs: rather than give both hawks in a hawk-hawk confrontation the *expected* payoff of -25, one is randomly chosen to receive a payoff of -100 while the other receives 50; similarly, when two doves meet, one is randomly chosen to receive a payoff of 40 while the other receives -10, instead of the expected payoff of 15. The payoffs of hawk-dove encounters are unchanged. Thus, after all pair-wise encounters occur, the two *strategies* still obtain equal fitnesses at the ESS; but, individual *agents* may do better or worse than others. We strongly suspect that the population of 6000 converges to all hawks because the higher number of encounters per agent allows expected payoff values to emerge. With a good enough approximation of expected payoff, the dynamics of the system are similar to noiseless payoffs, above.

We show elsewhere [2] that truncation selection is incapable of maintaining arbitrary fixed points. Indeed, with an *infinite* population playing the Hawk-Dove game ($w_0 = 26$, noiseless payoffs), truncation selection produces a fixed point of all hawks, chaos, or limit cycles, depending on the selection pressure. Given that truncation selection so completely disrupts the dynamics of the Hawk-Dove game with an infinite population, we believe the experiments in [3, 4, 5] that use truncation implicate the selection method more than other factors; the effects of noisy payoffs, incomplete mixing, and finite populations are inadequately isolated to warrant the stated conclusions.

3.2 Roulette Selection

The simulations of Fogel, et al, that do not use truncation selection use fitness-proportionate roulette-wheel selection instead [4]. Their core concern is whether evolutionary game theory can be applied to biological field study. Thus, if evolutionary game theory is applied to a finite population of physically embodied agents, such as animals, then clearly no pair-wise en-

counter will match an individual against itself. In this case, computation of fitness is slightly different than in Equation 2 because the “self” must be subtracted from the counts. For the Hawk-Dove game, we have:

$$f_H = E(H, H) * (p_H - \frac{1}{n}) + E(H, D) * p_D + w_0 \quad (6)$$

$$f_D = E(D, H) * p_H + E(D, D) * (p_D - \frac{1}{n}) + w_0$$

where f_i is the fitness of strategy i , $E(i, j)$ is the payoff for strategy i against strategy j , p_i is the proportion of strategy i in the population, and n is the population size.

Fogel, et al [4], recognize that variations in population size cause Equation 6 to change the proportion at which the two strategies reach equal fitness. In the limit of an infinite population, Equation 6 converges to the familiar ESS proportion of $\frac{7}{12}$ Hawks. For finite population sizes, however, the fitness-equilibrium proportion is actually higher. Their experiments that use fitness-proportional roulette-wheel selection give data that deviate with statistical significance from the theoretical ESS. But, the null hypothesis they choose to measure against is the equilibrium for an infinite population, and not that predicted by Equation 6, which varies according to population size. (Note that having an equilibrium higher than $\frac{7}{12}$ does not invalidate the above argument regarding truncation selection.)

Riley [9] raises a much more serious concern: Equation 6 can actually allow a strategy to meet Maynard-Smith’s criteria for an ESS (Equations 4 and 5), yet still be invadable. Such a strategy is termed a *quasi-equilibrium strategy*. Indeed, such is the case in the Hawk-Dove game. Thus, Equation 6 not only causes the equilibrium to diverge from the expected ESS, but also causes the ESS, in effect, to no longer exist. Riley goes on to show that, by introducing informational asymmetry into a game, a true (uninvadable) ESS is made to exist. The simulations of Fogel, et al, do not include informational asymmetry. Thus, divergence from the expected ESS, as observed in the simulations of [4], has ample analytical explanation quite apart from causes due to quantization noise or stochastic sampling, which are the factors that Fogel, et al, claim are implicated.

4 The Effects of Finite Populations: A Second Look

4.1 Assumptions

We now proceed with our own experiments. Since truncation selection appears pathological in the context of evolutionary game theory, we will concentrate

on proportional roulette-wheel selection in our simulations. We assume complete mixing and noiseless payoffs. We will also assume that an agent *can* play against itself. While this may lack biological plausibility, it is entirely possible—indeed common—in *co-evolutionary algorithms*. The important mathematical consequence is that we can revert back to Equation 2 and avoid the problems caused by Equation 6, above. Thus, we can better isolate the effects of noise due to quantization and stochastic sampling that arise with a finite population. Rather than use Maynard-Smith’s static conception of an ESS, we choose instead to use the formalism of dynamical systems to describe stability concepts. As we discuss below, the ESS of the Hawk-Dove game is an *attractive fixed point*.

4.2 Methods

We implement proportional roulette-wheel selection as follows: the fitness scores, f_i , are normalized, such that they sum to 1.0; a vector of sum *prefixes* is computed, such that the prefix value for agent i is the sum:

$$r_i = \sum_{k=1}^i f_k \quad (7)$$

We draw a value, x , from a uniform distribution and select the first agent whose prefix value is $\geq x$ to produce one offspring. This step is performed n times, where n is the size of the desired population.

We begin each experiment with the population at the theoretical ESS ratio of $p_H = \frac{7}{12}$ (and $p_D = \frac{5}{12}$). The population sizes for all experiments are chosen such that the ESS ratio is precisely representable by whole numbers of agents. Five different population sizes are tested: 60, 120, 300, 600, and 900. Each simulation is run for 2000 generations. The mean number of hawks in a run constitutes a single data point in an experiment; each experiment is repeated 100 times.

4.3 Results

Our results are shown in Figure 1. The x and y axes indicate the population size used in an experiment and the proportion of hawks in the population, respectively. Each circle represents the mean proportion of hawks seen during a single run. The dashed line indicates the ESS that evolutionary game theory predicts for an infinite population. The solid curve indicates the mean value of the 100 trials in each experiment. We see that for all population sizes, the mean value of hawks over all trials is consistently lower than the theoretical ESS proportion. As the population size

increases, the mean value asymptotically approaches the theoretical ESS. The second column of Table 1 lists the observed means and the fifth column gives t -values obtained when applying the t -test to the data from each experiment with respect to a null hypothesis $H_0 = .58333\dots$ (the ESS). In all cases, the observed data deviate with statistical significance from the theoretical ESS.




Figure 1: Data from finite-population simulations of Hawk-Dove game. Five experiments are conducted, each repeated 100 times. Each circle represents the mean proportion of hawks observed during a single 2000-generation simulation. The dashed line is the theoretical ESS. The solid curve indicates the mean values of the 100 data points in each experiment.

5 Analysis of Roulette Selection

That the data means come closer to the theoretical ESS as population size increases is perhaps easy to accept. What is less intuitive, however, is why the approach should be asymptotic from below the ESS value. To answer this question, we turn to dynamical systems theory [8, 11]. Figure 2 shows the *map diagram* for the Hawk-Dove game. Given a proportion of hawks, represented by the x -axis, one iteration of the evolutionary difference equation (Equation 3) will produce a new proportion, represented by the y -axis. The curve depicts the function that maps the proportion of hawks from one generation to the next. Intersections of the curve with the diagonal line indicate *fixed points*. If the slope of the curve at an intersection 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*. The Hawk-Dove game has

three fixed points: two unstable, where the population is either all hawks or all doves, and one stable, the polymorphic evolutionary stable strategy.

To see the dynamics of the Hawk-Dove game, one simply picks an initial point on the x-axis, draws a vertical line to the curve, and then alternately draws a horizontal line to the diagonal and then a vertical line to the curve until the ESS is reached. This procedure produces a *cobweb diagram* that indicates the *orbit* of the initial point. Two such orbits are shown in Figure 2, which begin ± 0.25 away from the ESS ($p_H = \frac{7}{12}$). The key observation is that the orbit that begins below the ESS requires more iterations to reach the ESS than does the orbit that begins above. Figure 3 shows the number of iterations required for every initial condition from 0 to 1, at intervals of 0.001, to arrive within $\epsilon = 0.0001$ of the ESS. For any δ , orbits starting at $\frac{7}{12} - \delta$ require \geq iterations than orbits starting at $\frac{7}{12} + \delta$.

This observation is important because the operation of the roulette-wheel method of proportional selection produces a *binomial distribution* around the desired proportion of hawks indicated in the wheel. The binomial distribution obtained for n spins of a roulette wheel that gives a hawk with probability p is:

$$\text{bin}(n, p) = \binom{n}{i} * p^i * (1 - p)^{n-i}, \text{ for } i = 0 \dots n \quad (8)$$

That is, the probability of having exactly i hawks (and therefore $n - i$ doves) in the next generation of n agents is the probability of getting i hawks times the probability of getting $n - i$ doves times the number of ways n spins can result in i hawks. The resulting distribution for all possible outcomes, from no hawks to no doves, has a mean value of $p * n$ hawks—the desired proportion indicated in the wheel.

Figure 4 shows the “spread” of the binomial distribution for values of $n = \{60, 120, 300, 600, 900\}$ (the population sizes used in our simulations, above) and $p = \frac{7}{12}$ (the desired proportion of hawks at the ESS). Each curve in Figure 4 represents the possible outcomes that sum to 99 percent of likelihood. The spread clearly widens as the population size decreases (and the law of large numbers exerts less influence).

Returning to Figure 3, we see that an increase in spread at the ESS will cause the asymmetry in convergence time to grow. This is what causes the observed population averages to tend to fall below the theoretical ESS by an amount inversely related to population size. But, how might we approximate the actual divergence without resorting to empirical methods? We offer an approach based on Markov chains.




Figure 2: Map diagram of Hawk-Dove game. The curve depicts the function that maps the proportion of hawks from one generation to the next.




Figure 3: Number of iterations needed to arrive within $\epsilon = 0.0001$ of the ESS.




Figure 4: Spread of binomial distribution at ESS for various population sizes.

6 Calculating the Divergence

6.1 Method

In this section, we develop the intuition from above to predict how experimental data will deviate from the theoretical ESS value. For the Hawk-Dove game (or any two-strategy game), there are only $n + 1$ possible states in which a population of size n can be, from 0 hawks to n hawks. These can be considered the states of a Markov chain. At each population state, i , the binomial distribution of the roulette wheel gives a set of *transition probabilities* to states $j = 0 \dots n$. This information is represented by a *transition matrix*, M :

$$M = \begin{bmatrix} \Pr(0|0) & \dots & \Pr(0|n) \\ \vdots & \ddots & \vdots \\ \Pr(n|0) & \dots & \Pr(n|n) \end{bmatrix} \quad (9)$$

where each column gives the transition probabilities, $p(j|i)$, from state i to states $j = 0 \dots n$. Each column sums to 1.0.

To compute the matrix columns, we need only calculate the appropriate binomial distribution for each possible state of the population. Recall that for any *desired* proportion of hawks, d , the binomial distribution $\text{bin}(n, d)$ tells us the likelihoods of all *possible* proportions of hawks, given n spins of the roulette wheel. For a population with i hawks (a proportion of i/n), the desired number of hawks for the next generation is $d = \text{map}(i/n)$, where the function *map* corresponds to an iteration of the map diagram in Figure 2. Thus, the columns of the transition matrix, M , are:

$$M = [\text{bin}(n, \text{map}(\frac{0}{n})) \dots \text{bin}(n, \text{map}(\frac{n}{n}))] \quad (10)$$

Let us consider an initial population composed of exactly $\frac{7}{12}$ hawks, and represent this population as a column vector, b , of probabilities on the different possible population states:

$$b = \begin{bmatrix} \Pr(0) \\ \vdots \\ \Pr(\frac{7}{12}n - 1) \\ \Pr(\frac{7}{12}n) \\ \Pr(\frac{7}{12}n + 1) \\ \vdots \\ \Pr(n) \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad (11)$$

The distribution of possible populations after t gener-

ations is M raised to power t times b :

$$b^t = (M)^t * b \quad (12)$$

Therefore, the expected proportion of hawks after t generations is given by the weighted sum:

$$E[\%hawks]^t = \sum_{i=0}^n \frac{b_i^t * i}{n} \quad (13)$$

What is the limit behavior, as time goes to infinity? Observe that roulette selection always gives a non-zero probability of arriving at a state of all hawks or all doves. And, since a population of all hawks will remain all hawks, and a population of all doves will remain all doves, we know that, in infinite time, one of these two *absorbing states* is an inevitable outcome. The probabilities of these two absorbing states are sensitive to the initial state.

Thus, the distribution, b^∞ must have the form:

$$b^\infty = [\Pr(0); 0; 0; \dots; 0; 0; \Pr(n)] \quad (14)$$

where only states 0 and n have non-zero probability.

Referring back to Equation 13, this implies that the expected proportion of hawks is simply the probability of ending in the absorbing state of all hawks. Thus, $E[\%hawks]^\infty = b_n^\infty$.

6.2 Predictions and Data

For the population sizes used in our experiments, the expected number of hawks for even small values of t gives an excellent approximation of the limit behavior: $|b_n^\infty - b_n^t| < \epsilon$ for small t . Figure 5 shows the values of Equation 13 for values of $t = 1 \dots 30$. We see that the limiting distribution is approached asymptotically. Our approximate expected values (computed at $t = 100$) are listed in column three of Table 1. We see that they are very close, indeed, to the actual means observed in our experiments (shown in column two). Further, using these predictions as null hypotheses in the t -test, the resulting t -values (shown in column 4) indicate that none of the predictions can be rejected (95% confidence level). Thus, the dynamics and equilibria of evolutionary game theory have been demonstrated to exist and apply to finite populations.

7 Stochastic Universal Sampling

While the predicted divergence from the theoretical ESS is clear when we look at an ensemble of 100 data points, an individual run can show a rather large

Table 1: Simulation results and predictions.

Size	Mean Hawks	Adj. ESS	$t_{Adj.}$	t_{ESS}
60	0.578166	0.578008	0.4922	-16.09
120	0.580682	0.580755	-0.3289	-11.99
300	0.582301	0.582320	-0.1267	-6.72
600	0.582864	0.582829	0.3953	-5.31
900	0.582981	0.582998	-0.2132	-4.47




Figure 5: Adjustment of ESS. Though small in magnitude, the adjusted values vary from the ESS enough to cause statistically significant deviation from $\frac{7}{12}$ in simulation data (indicated by column 5 of Table 1).

amount of noise due to the roulette wheel. Figure 6 shows 500 generations of the Hawk-Dove game beginning at the ESS for a population of size 60. Though noise is reduced with a larger population, there exists an alternative selection scheme that reduces the noise even further while maintaining a small population.

Baker’s *Stochastic Universal Sampling* (SUS) [1] provides fitness-proportionate selection with minimal use of a stochastic process. Simply, rather than spin a roulette wheel with one “pointer” on it n times, we spin a roulette wheel with n equally-spaced pointers just once. This method guarantees that an individual that should appear with proportion, p_{Ideal} , will appear with proportion $\lfloor p_{Ideal} * n \rfloor / n \leq p_{Actual} \leq \lceil p_{Ideal} * n \rceil / n$. Baker’s analysis of this method indicates that it has excellent statistical properties.

Indeed, SUS produces far less noisy results than ordinary proportional roulette selection, as seen in Figure 7. Because the ESS can be exactly represented by a population of size 60, we instead choose a population of 61. We see that the actual proportions remain as close to the theoretical ESS as the resolution of the

population allows. In fact, SUS will converge to the ESS within the resolution of the population regardless of the initial condition ($0 < p_H^{Initial} < 1.0$). For a population of size 100, SUS selection gives data that cannot be rejected by the t -test (at 95% confidence level) with a null hypothesis of $H_0 = \frac{7}{12}$. That is, because SUS uses only a single “spin,” the distribution of possible outcomes is made so narrow that the calculation of divergence from the theoretical ESS is unnecessary.




Figure 6: Sample run with population size 60 using proportional roulette-wheel selection.




Figure 7: Sample run with population size 61 using Baker’s *Stochastic Universal Sampling* selection [1].

8 Conclusion

There is no question that the introduction of noise, be it from the finiteness of a population, variation in pay-

offs, or incomplete mixing, affects the dynamics and equilibria in evolutionary game theory. And, Fogel, et al, are correct to highlight the issue. But, the selection method distorts the process, as well—it can never be made transparent. In simulations that deal with finite populations, reproduction generally cannot occur in exact proportion to fitness—agents can only replicate whole-numbers of offspring. Thus, the question of how one is to *implement* reproduction, fitness-proportionate or otherwise, arises. All these sources of “noise” determine whether the essential character of evolutionary game theory remains intact.

Truncation selection is shown to be pathological in the context of evolutionary game theory—it can neither attain nor maintain a polymorphic ESS (even with infinite populations). Roulette-wheel selection is shown to diverge from the theoretical ESS due to the interaction between the wheel’s binomial distribution and the convergence properties of the Hawk-Dove map (Figure 2). We show how the expected divergence can be computed without empirical data. Finally, we show that Baker’s SUS selection method allows a finite population of modest size to approximate the ESS without statistically significant deviation. If the population size n can be divided into whole numbers that precisely represent the ESS ratio, then SUS selection allows the population to converge to the exact ESS. Thus, we have demonstrated that the dynamics and equilibria of evolutionary game theory can persist with finite populations, provided that the selection method is appropriately chosen and implemented.

Of course, this study has only addressed the smallest of games; what happens when the number of strategies in a population is comparable to the population’s size? This is akin to representing the state of a dynamical system with very low resolution, and we will need to return to dynamical systems theory for a satisfactory answer.

While the primary concern of Fogel, et al, is the applicability of evolutionary game theory to biological study, we make no attempt to relate our results and methods to a biological context—to avoid the problems and remedies discussed by Riley [9], we assume an agent can play against itself. Nevertheless, we believe our approach suggests a novel framework in which to understand the dynamics of coevolutionary algorithms—for evolutionary game theory can be understood as a coevolutionary algorithm without variational operators. This avenue is another subject of our current research.

9 Acknowledgments

The authors wish to thank members of the DEMO Lab and particularly Ofer Melnik for many fruitful discussions and suggestions.

References

- [1] J. E. Baker. Reducing bias and inefficiency in the selection algorithm. In *Proceedings of the second international conference on Genetic Algorithms*, pages 14–21. Lawrence Erlbaum, 1987.
- [2] S. G. Ficici, O. Melnik, and J. B. Pollack. A game-theoretic investigation of selection methods used in evolutionary algorithms. In A. Zalzal et al., editors, *2000 Congress on Evolutionary Computation*. IEEE Press, 2000.
- [3] D. B. Fogel and G. B. Fogel. Evolutionary stable strategies are not always stable under evolutionary dynamics. In *Evolutionary Programming IV*, pages 565–577, 1995.
- [4] D. B. Fogel, G. B. Fogel, and P. C. Andrews. On the instability of evolutionary stable strategies. *BioSystems*, 44:135–152, 1997.
- [5] G. B. Fogel, P. C. Andrews, and D. B. Fogel. On the instability of evolutionary stable strategies in small populations. *Ecological Modelling*, 109:283–294, 1998.
- [6] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.
- [7] J. Maynard-Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982.
- [8] H.-O. Peitgen, H. Jürgens, and D. Saupe. *Chaos and Fractals: New Frontiers of Science*. Springer-Verlag, 1992.
- [9] J. G. Riley. Evolutionary equilibrium strategies. *Journal of Theoretical Biology*, 76:109–123, 1979.
- [10] L. Samuelson. *Evolutionary Games and Equilibrium Selection*. MIT Press, 1997.
- [11] S. H. Strogatz. *Nonlinear Dynamics and Chaos*. Addison Wesley, 1994.
- [12] J. Weibull. *Evolutionary Game Theory*. MIT Press, 1995.