Evolutionary Dynamics of Finite Populations in Games with Polymorphic Fitness-Equilibria

Sevan G. Ficici *, Jordan B. Pollack

Department of Computer Science, Brandeis University, Waltham MA 02454 USA

Abstract

The Hawk-Dove (HD) game, as defined by Maynard Smith (1982), allows for a polymorphic fitness-equilibrium (PFE) to exist between its two pure strategies; this polymorphism is the attractor of the standard replicator dynamics (Taylor and Jonker, 1978; Hofbauer and Sigmund, 1998) operating on an infinite population of pure-strategists. Here, we consider stochastic replicator dynamics, operating on a finite population of pure-strategists playing games similar to HD; in particular, we examine the transient behavior of the system, before it enters an absorbing state due to sampling error. Though stochastic replication prevents the population from fixing onto the PFE, selection always favors the under-represented strategy. Thus, we may naively expect that the mean population state (of the pre-absorption transient) will correspond to the PFE. The empirical results of Fogel et al. (1997) show that the mean population state, in fact, deviates from the PFE with statistical significance. We provide theoretical results that explain their observations. We show that such deviation away from the PFE occurs when the selection pressures that surround the fitness-equilibrium point are asymmetric. Further, we analyze a Markov model to prove that a finite population will generate a distribution over population states that equilibrates selection-pressure asymmetry; the mean of this distribution is generally not the fitness-equilibrium state.

Preprint submitted to Journal of Theoretical Biology *2 March 2007*

Key words: Finite population, Replicator dynamics, Selection pressure, Fitness equilibrium

1 Introduction

The standard replicator dynamics (e.g., Taylor and Jonker, 1978; Hofbauer and Sigmund, 1998) are deterministic processes that operate on infinite populations. Here we examine two stochastic replicator dynamics that operate on small, well-mixed finite populations of fixed-size N. These two replicator dynamics can be described as frequency-dependent Wright-Fisher (Wright, 1931; Fisher, 1930) and Moran (Moran, 1958) processes, respectively, operating on haploid populations. A population consists of pure-strategist agents; these agents play a symmetric 2x2 variable-sum game. We are particularly interested in games, such as Hawk-Dove (HD) (Maynard Smith, 1982), that have a polymorphic fitness-equilibrium (PFE).

In the games we study, when the population state is away from polymorphic fitness-equilibrium, then the under-represented strategy (relative to the strategy proportions at the PFE) is always favored by selection. Thus, selection always acts to move the population state towards the PFE. Given standard replicator dynamics operating deterministically on an infinite population, this action of selection causes the PFE to be a point-attractor (Hofbauer and Sig-

* Corresponding author. Present address: Maxwell-Dworkin Laboratory, Division of Engineering and Applied Sciences, Harvard University. 33 Oxford Street, Room 242, Cambridge, Massachusetts, 02138 USA; email: sevan@eecs.harvard.edu; Phone +1.617.495.9289; Fax +1.617.496.1066.

Email address: pollack@cs.brandeis.edu (Jordan B. Pollack).

mund, 1998). Given a stochastic finite-population system, however, the role of the PFE in the system's behavior is less clear. Indeed, if our system lacks mutation, then we know that sampling error will ultimately cause the population to enter one of two monomorphic absorbing states; but, the pre-absorption transient can be very long-lived—how does the PFE shape the dynamics of the population before absorption occurs? If, instead, the system includes mutation, then it will have a unique steady-state distribution over the possible population states; does the PFE correspond to the expected population-state at the system's steady-state?

In an empirical investigation of finite-population dynamics using agent-based computer simulation, Fogel et al. (1997) observe that the mean population-state obtained under a stochastic replication process diverges from the PFE with statistical significance. We provide theoretical explication of this observation. We show that deviation away from the PFE occurs when the selection pressures that surround the PFE are asymmetric. Game payoffs determine the magnitude and shape of this asymmetry, but the amount of asymmetry to which the system is actually exposed is determined by the population's size; smaller populations are more exposed to selection-pressure asymmetry and so diverge more from the PFE.

Further, we prove with Markov-chain analysis that the finite-population process generates a distribution over population states that equilibrates asymmetries in selection pressure; the mean of the distribution is generally not the PFE. More simply put, the finite populations we study equilibrate selection pressure, not fitness.

This article is organized as follows. Section 2 reviews related work. Section

3 details the finite-population models we examine; this section specifies the class of games that we consider, discusses the calculation of the PFE, and describes the replicator dynamics that we analyze. Section 4 presents four example games that we examine in detail. Section 5 gives empirical results on these games which indicate that finite-population behavior generally does not correspond to the PFE. Section 6 proposes the hypothesis that asymmetry in selection pressure causes actual behavior to diverge from the PFE, and Section 7 formalizes this intuition. Sections 8 and 9 provide further discussion and concluding remarks. Appendix A details our agent-based simulation methods. Appendix B describes how we construct our Markov models, and Appendix C gives our central proof. Appendix D discusses the special case of very small populations in the absence of mutation. Appendix E contrasts our equation to compute fitness equilibrium in a finite population with the equation given by Schaffer (1988) to compute an evolutionarily stable strategy (ESS) in a finite population.

2 Related Work

Our focus on fixed-size populations of pure-strategists and polymorphic fitness-equilibria stands in contrast to much other research in finite-population dynamics. For example, in their studies of evolutionarily stable strategies, Schaffer (1988), Vickery (1987, 1988), and Maynard Smith (1988) require agents to use mixed-strategies. Schaffer (1988) points out that a population of N pure-strategists cannot represent an arbitrary distribution over a game's pure strategies. This is certainly true for a static population; but, when acted upon over time by a stochastic process, we can discuss the population's expected

state, which can in general be an arbitrary distribution. Thus, while stochastic replication prevents the population from converging onto the PFE, we may naively suppose that the PFE will accurately describe the expected population state.

More recent work by Bergstrom and Godfrey-Smith (1998) and Orzack and Hines (2005) expands upon these earlier studies to better determine the circumstances under which a mixed-strategy ESS will arise; Orzack and Hines (2005), in particular, examine a wide range of initial conditions, and consider the relationship between population size, genetic drift, and the strength of selection (we have more to say about these relationships, below). Other work on finite populations attends to issues such as invadability (Riley, 1979) and fixation probabilities of mutants (Nowak et al., 2004; Taylor et al., 2004; Lessard, 2005). Schreiber (2001) and Benaïm et al. (2004) consider the case where the population size is allowed to grow and shrink.

Most relevant to our work are the empirical investigations of Fogel et al. (1997); using agent-based computer simulation, Fogel et al. (1997) examine the dynamics of a population of 100 pure-strategists playing the Hawk-Dove game. In addition to using stochastic fitness-proportional replication processes, Fogel et al. (1997) consider finite-population dynamics under stochastic (and possibly incomplete) mixing and stochastic payoffs (instead of expected values). Their empirical results show that the mean population-state deviates with statistical significance from the game's PFE. (Fogel and Fogel (1995) and Fogel et al. (1998), as well as other experiments in Fogel et al. (1997), use truncation-like replication processes with finite populations; see Ficici et al. (2000, 2005) for further analyses on truncation dynamics.) The work we present in this article focuses on stochastic selection and so uses (deterministic) complete mixing

and expected payoffs.

Liekens et al. (2004) adopt the general methodology of our earlier work on finite populations (Ficici and Pollack, 2000) and corroborate our early results, but do so in a system that extends Ficici and Pollack (2000) to include mutation; this modification allows the system to be modeled by an irreducible Markov chain that has a unique steady-state distribution. For the present study, we begin with a zero-mutation system similar to Ficici and Pollack (2000); we model this system with a reducible Markov chain. When we arrive to our proof, we move to an irreducible Markov chain, which allows our results to generalize to the case where mutation is used.

This article expands our initial research on finite populations (Ficici and Pollack, 2000). While our original work suggests that selection-pressure asymmetry causes the observed deviation away from the PFE, it lacks the formal argument given here. Our original work limits itself to the Wright-Fisher process and assumes self-play; we now also examine the Moran process as well as the case where agents cannot interact with themselves.

3 Finite-Population Models

3.1 2x2 Symmetric Games

A generic payoff matrix for a symmetric two-player game of two pure strategies X and Y is given by,

$$\mathbf{G} = \begin{array}{c} \\ X \\ \\ Y \end{array} \begin{array}{c} \begin{array}{cc} X & Y \end{array} \\ \begin{bmatrix} a & b \\ \\ c & d \end{bmatrix} \end{array}.$$

By convention, the payoffs are for the row player.

There exist exactly two payoff structures for such a game that create a polymorphic fitness-equilibrium between X- and Y-strategists. Case 1 has $a > c$ and $b < d$. In games with this payoff structure, selection pressure always points away from the polymorphic fitness-equilibrium, making it dynamically unstable. Case 2 has $a < c$ and $b > d$. As we discuss below, games with this second payoff structure have selection pressure always pointing towards the polymorphic fitness-equilibrium, which creates a negative feedback-loop around the polymorphism. In this article, we are interested only in Case 2.

3.2 Cumulative Rewards in Infinite Populations

Infinite-population models typically assume *complete mixing* (Maynard Smith, 1982); that is, each agent interacts with every other agent in the population, accumulating payoff as it goes. Let $p \in [0, 1]$ be the proportion of X-strategists in the population, with $(1 - p)$ being the proportion of Y-strategists. The cumulative rewards obtained by X- and Y-strategists, denoted w_{X} and w_{Y}, respectively, are linear functions of p, where

$$w_{\mathrm{X}} = pa + (1-p)b + w_0,$$
$$w_{\mathrm{Y}} = pc + (1-p)d + w_0. \tag{1}$$

The constant w_0 is a "background fitness" that is large enough to ensure that w_{X} and w_{Y} are non-negative (in the discrete-time replicator dynamics, which we use below, agents reproduce offspring in proportion to cumulative reward, and an agent cannot have less than zero offspring). Note that for every pair $\langle \mathbf{G}, w_0 \rangle$ there exists the pair $\langle \mathbf{G}', w_0' \rangle$, where $\mathbf{G}' = \mathbf{G} + w_0$ and $w_0' = 0$, that gives identical replicator behavior in all respects; this is a many-to-one mapping. This equivalence is easy to check algebraically in (1).

Given the payoff structure $a < c$ and $b > d$ (Case 2, above), we can easily see that a polymorphic fitness-equilibrium must exist, and that selection always points to it. When $p = 1$, we have $w_{\mathrm{X}} = a < c = w_{\mathrm{Y}}$; at $p = 0$, we have $w_{\mathrm{X}} = b > d = w_{\mathrm{Y}}$. Thus, the lines described by w_{X} and w_{Y} must intersect at some value of p; the two strategies are at fitness equilibrium at this intersection. We denote the location of this intersection as p_{Eq_∞}. For $p < p_{\mathrm{Eq}_\infty}$, we have an overabundance of Y-strategists and selection favors X-strategists ($w_{\mathrm{X}} > w_{\mathrm{Y}}$); for $p > p_{\mathrm{Eq}_\infty}$, we have the opposite situation and selection favors Y-strategists ($w_{\mathrm{X}} < w_{\mathrm{Y}}$). This negative feedback-loop is responsible for making p_{Eq_∞} a stable fixed-point of the standard replicator dynamics operating on an infinite-population (we discuss standard discrete-time replicators below). To calculate the location of the polymorphic fitness equilibrium, we simply set $w_{\mathrm{X}} = w_{\mathrm{Y}}$ and solve for p; doing so, we find the polymorphic fitness equilibrium p_{Eq_∞} to be

$$p_{\mathrm{Eq}_\infty} = \frac{d-b}{a-b-c+d}. \tag{2}$$

Note that p_{Eq_∞} is independent of w_0.

3.3 Cumulative Rewards in Finite Populations

If we have a finite population of size N, we must consider whether an agent is allowed to interact with itself or not. If we allow self-play, then we continue to use (1) to calculate cumulative reward; this means that the location of the polymorphic fitness-equilibrium is unchanged. If we assume that an agent cannot interact with itself, then the numbers of X- and Y-strategists that an agent sees depends upon the identity of the agent. Cumulative rewards in a finite population without self-play are given by

$$\begin{aligned} w_{\mathrm{X}} &= \frac{pN-1}{N-1}a + \frac{(1-p)N}{N-1}b + w_0, \\ w_{\mathrm{Y}} &= \frac{pN}{N-1}c + \frac{(1-p)N-1}{N-1}d + w_0. \end{aligned} \tag{3}$$

Though our equations for cumulative rewards have changed to exclude self-play, they remain linear in p and selection continues to point towards the fitness equilibrium. As with (1), for every pair $\langle \mathbf{G}, w_0 \rangle$ there exists the pair $\langle \mathbf{G}', w_0' \rangle$, where $\mathbf{G}' = \mathbf{G} + w_0$ and $w_0' = 0$, that gives identical replicator behavior in all respects.

We can easily see that, as $N \to \infty$, (3) converges to the infinite-population (or self-play) rewards given by (1). Thus, any replicator dynamic built upon

(1) can be approximated with arbitrary fidelity by the same dynamics built upon (3), given a sufficiently large population.

Let p_{Eq_N} denote the location of the polymorphic fitness-equilibrium in a finite population. To calculate p_{Eq_N}, we again set $w_{\mathrm{X}} = w_{\mathrm{Y}}$ and solve for p; the polymorphic fitness-equilibrium in a finite population without self-play is thus

$$p_{\mathrm{Eq}_N} = \frac{\frac{a-d}{N} + d - b}{a - b - c + d}. \tag{4}$$

We again notice that the equilibrium is independent of w_0, but it does now generally depend upon the population size. An exception occurs when payoffs a and d are equal; in this case, $p_{\mathrm{Eq}_N} = p_{\mathrm{Eq}_\infty}$ regardless of population size. Equation (4) converges to (2) as $N \to \infty$. If $a < d$, then p_{Eq_N} asymptotically approaches p_{Eq_∞} from above (the payoff structure of Case 2 implies that the denominator of (4) is negative); if $a > d$, then p_{Eq_∞} is approached from below. Appendix E contrasts (4) with the equation given by Schaffer (1988) to compute an evolutionarily stable strategy in a finite population.

3.4 Wright-Fisher Replicator

The generic discrete-time Wright-Fisher replicator equation (Hofbauer and Sigmund, 1998; Wright, 1931; Fisher, 1930) for a 2x2 game is given by

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

Instantiating w_{X} and w_{Y} in (5) with (1), we obtain

$$f(p) = \frac{p^2(a-b) + pb + pw_0}{p^2(a-b-c+d) + p(b+c-2d) + d + w_0}, \tag{6}$$

which can be interpreted as a deterministic mapping, given an infinite population, from population state p to state $f(p)$; this mapping describes the exact behavior of the population. Alternatively, we may interpret (6) as a mean-field equation describing the expected transformation, under self play, of a finite population in state p. If we instantiate w_{X} and w_{Y} with (3), then we obtain

$$f(p) = \frac{\Big(p^2(a-b) + pb + pw_0\Big)N - pa - pw_0}{\Big(p^2(a-b-c+d) + p(b+c-2d) + d + w_0\Big)N + p(d-a) - d - w_0}, \tag{7}$$

which represents the mean-field equation for a finite population without self-play. Note that (7) converges onto (6) as $N \to \infty$.

Unlike the fitness equilibrium equations (2) and (4), the replicator equations (6) and (7) depend upon w_0. Larger values of w_0 decrease the magnitude of fitness differential between w_{X} and w_{Y}, as is clear from (1) and (3). As a result, larger w_0 weaken selection pressure and reduce the expected change in population state.

In the games we study, the polymorphism p_{Eq_∞} is the only stable fixed-point of the infinite-population replicator dynamics. Given a finite population, the replicator dynamics operate stochastically and the mapping from p to $f(p)$ describes only the expected transformation of the population state. Consequently, though selection always favors moving p towards p_{Eq_N}, the population cannot fix onto this polymorphic fitness-equilibrium because replication

is noisy.

The Wright-Fisher process is generational—each iteration creates a new population of N agents. We take N i.i.d. samples from a random variable $\mathcal{X}$. Given a population with state p, the probability of creating an X-strategist offspring on a particular sample is $\Pr(\mathcal{X} = X|p) = f(p)$. This creates a binomial distribution with expected value $f(p)$.

In the absence of mutation, stochastic replication will eventually cause the population to enter one of two monomorphic absorbing states (all X- or all Y-strategists). The expected time to absorption can be extremely long, however, which invites investigation of the pre-absorption transient. We show below that the PFE poorly predicts the mean population-state of the pre-absorption transient in our stochastic process. The purpose of this article is to elucidate why this is so and explain how the polymorphic fitness-equilibrium relates to the mean population-state.

3.5 Moran Replicator

In contrast to the Wright-Fisher model, the Moran process (Moran, 1958) creates only a single offspring in each iteration; in this article, we study a version of the Moran process that includes frequency-dependent fitness, such as that used in Nowak et al. (2004). The process involves two stochastic steps: first we select (with probability determined by cumulative payoff) an agent to parent one offspring, and then we select (with uniform probability) an agent from the population for replacement (possibly the parent agent) and insert the offspring.

The generic mean-field equation for the Moran process is given by

$$f(p) = \Pr(X^0)p + \Pr(X^+)(p + \frac{1}{N}) + \Pr(X^-)(p - \frac{1}{N}), \qquad (8)$$

where the probability $\Pr(X^0)$ that the offspring plays the same strategy as the agent it replaces (resulting in no change of population state) is

$$\Pr(X^0) = 1 - \Pr(X^+) - \Pr(X^-),$$

the probability $\Pr(X^+)$ that an X-strategist offspring replaces a Y-strategist (resulting in an increase of p) is

$$\Pr(X^+) = \Pr(\mathcal{X} = X|p)(1 - p),$$

the probability $\Pr(X^-)$ that a Y-strategist offspring replaces an X-strategist (resulting in a decrease in p) is

$$\Pr(X^-) = \Pr(\mathcal{X} = Y|p)p,$$

and the probabilities $\Pr(\mathcal{X} = X|p)$ and $\Pr(\mathcal{X} = Y|p)$ of creating an X- and Y-strategist offspring given the population state p, respectively, are

$$\Pr(\mathcal{X} = X|p) = \frac{pw_{\mathrm{X}}}{pw_{\mathrm{X}} + (1-p)w_{\mathrm{Y}}},$$

$$\Pr(\mathcal{X} = Y|p) = \frac{(1-p)w_{\mathrm{Y}}}{pw_{\mathrm{X}} + (1-p)w_{\mathrm{Y}}}.$$

Because we only replace a single agent per iteration, $\lim_{N\to\infty} f(p) = p$. Since the *effective* population size of the Moran process is $N/2$ (Ewens, 2004; Wakeley and Takahashi, 2004), we can approximate a new generation with only $N/2$

iterations (nevertheless, we use N iterations in this article, which is harmless for our purposes).

When we instantiate the generic Moran equation (8) for the case of agent self-play (1), we get

$$f(p) = p - p/N + \frac{p^2(a-b) + p(b+w_0)}{\alpha N}, \tag{9}$$

$$\alpha = p^2(a-b-c+d) + p(b+c-2d) + d + w_0. \tag{10}$$

Instantiating the Moran equation (8) for the case of no self-play (3) gives

$$f(p) = \frac{pN^2\alpha - pN\Big(\alpha + p(b-d) + d - b\Big) + (p-p^2)(d-a)}{N^2\alpha + N\Big(p(d-a) - d - w_0\Big)}, \tag{11}$$

where α is given by (10).

4 Example Games

We will examine the following four symmetric 2x2 games (we use $w_0 = 0$ for all four games):

$$\mathbf{G}_1 = \begin{bmatrix} 0.01 & 9.51 \\ 2.51 & 6.01 \end{bmatrix}, \quad \mathbf{G}_2 = \begin{bmatrix} 5.61 & 5.61 \\ \frac{6227}{700} & 1.01 \end{bmatrix},$$

$$\mathbf{G}_3 = \begin{bmatrix} 1 & 4 \\ \frac{22}{7} & 1 \end{bmatrix}, \qquad \mathbf{G}_4 = \begin{bmatrix} 0.01 & 10.01 \\ \frac{5007}{700} & 0.01 \end{bmatrix}.$$

Given a finite population with self-play (or an infinite population), all four games have polymorphic fitness-equilibrium at $p_{\mathrm{Eq}_N} = 7/12$. Given a finite population without self-play, we know from (4) that, as N increases, p_{Eq_N} will approach $7/12$ from above for $\mathbf{G}_1$ and from below for $\mathbf{G}_2$; games $\mathbf{G}_3$ and $\mathbf{G}_4$ have fitness-equilibrium at $7/12$ regardless of population size. Figure 1 plots p_{Eq_N}, calculated using (4), for each game over a range of population sizes.

The apparent precision with which game payoffs are expressed is deceptive and should not lead the reader to believe that we are describing singularities or special cases; we merely choose payoffs that elicit a nice variety of behaviors from our finite-population system. We will see below that our experimental results yield mean population-states that generally disagree with fitness-equilibrium; we will then go on to show how the variety of behaviors that emerge can be unified by understanding that our finite populations equilibrate selection pressure, not fitness.

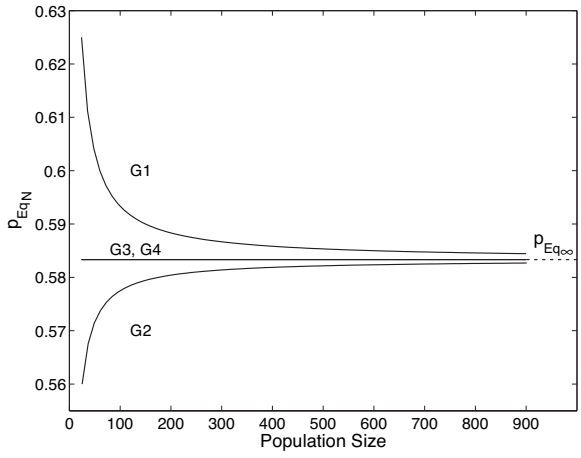

Fig. 1. Proportion p_{Eq_N} of X-strategists in population at polymorphic fitness-equilibrium (without self-play) for games $\mathbf{G}_1$ through $\mathbf{G}_4$ over a range of population sizes; values are calculated using (4). With self-play or an infinite population (indicated by dashed line), polymorphic fitness-equilibrium is $p_{\mathrm{Eq}_N} = 7/12$ for all four games.

5 Agent-Based Simulations and Markov-Chain Models

We now examine how well the polymorphic fitness-equilibrium curves in Figure 1 predict the behavior of stochastic agent-based simulations and Markov-chain models. We investigate Wright-Fisher and Moran dynamics, with and without agent self-play. We choose population sizes N such that $N \cdot p_{\mathrm{Eq}_N}$ is a whole number of agents. For all experiments, we initialize the population to be at polymorphic fitness-equilibrium.

We use 19 different population sizes in the range $N \in [24, 900]$; when played without self-play, game $\mathbf{G}_2$ requires an additional agent to obtain whole numbers at the PFE (i.e., $N \in [25, 901]$). (Note that the smallest population size used in this section for $\mathbf{G}_2$ under Moran dynamics without self-play is $N = 37$; we devote Appendix D to the case where $N = 25$.)

Our agent-based simulations, detailed in Appendix A, do not include mutation, and so we obtain a system with monomorphic absorbing states. The corresponding Markov-chain models, detailed in Appendix B, are thus reducible. The data we present in this section concern only the pre-absorption dynamics of the system. Let $\overline{p}$ denote the mean population-state of the pre-absorption transients observed in agent-based trials. Let $\mathrm{E}[p|t]$ denote the expected population-state at time t as determined by the Markov-chain model; let T_{A} denote the expected number of time-steps before absorption occurs, and $\mathrm{E}[p]$ denote the expected population state at time $t = T_{\mathrm{A}}$. We use the two-tailed t-test at significance-level 0.001 when comparing $\overline{p}$ to p_{Eq_N}.

5.1 Wright-Fisher Replication

The four panels of Figure 2 give Markov-chain and agent-based simulation results for our example games under the Wright-Fisher process without agent self-play. Results obtained from the Markov-chain model (indicated by $\mathrm{E}[p]$) agree well with empirical data collected from agent-based simulation (indicated by $\overline{p}$), but they generally do not match the polymorphic fitness-equilibrium curves (indicated by p_{Eq_N}). For game $\mathbf{G}_1$, the expected population-state $\mathrm{E}[p]$ is sandwiched between p_{Eq_N} and p_{Eq_∞}; the mean population-state $\overline{p}$ deviates from p_{Eq_N} with statistical significance. Thus, even though we initialize the population to be at polymorphic fitness-equilibrium, selection pressure somehow systematically moves the mean population state away from p_{Eq_N}.

In game $\mathbf{G}_2$, both $\mathrm{E}[p]$ and $\overline{p}$ are consistently above p_{Eq_N}, and $\overline{p}$ deviates with statistical significance from p_{Eq_N} (though $\overline{p}$ appears to correspond nicely with the infinite-population equilibrium p_{Eq_∞}). The slight non-monotonic progres-

sion in E[p] for $N \in \{25, 37, 49\}$, shown magnified in the inset, is a manifestation of very small populations (discussed in Appendix D).

Games $\mathbf{G}_3$ and $\mathbf{G}_4$ both have $p_{\mathrm{Eq}_N} = 7/12$ regardless of population size, yet they yield very different results. In $\mathbf{G}_3$, E[p] asymptotically approaches 7/12 from above; agent-based simulations with population-sizes $N \leq 300$ give results that deviate with statistical significance from fitness-equilibrium (though the magnitude of divergence is less than that in games $\mathbf{G}_1$ and $\mathbf{G}_2$); population sizes $N > 300$ (indicated on the graph by filled squares) do not deviate with statistical significance. In game $\mathbf{G}_4$, in contrast to the other three games, we cannot discern any statistically significant deviation from fitness-equilibrium for any population size; the only systematic change is the decrease in standard deviation as population size increases.

Figure 3 gives our results for the Wright-Fisher replicator when agent self-play is allowed. In this case, $p_{\mathrm{Eq}_N} = p_{\mathrm{Eq}_\infty} = 7/12$ for all games over all population sizes. The gaps between E[p] and p_{Eq_N} are slightly, but consistently, larger with self-play than without self-play. Again, game $\mathbf{G}_1$ shows E[p] $< p_{\mathrm{Eq}_N}$, while $\mathbf{G}_2$ gives E[p] $> p_{\mathrm{Eq}_N}$; in both games, $\overline{p}$ deviates with statistical significance from p_{Eq_N} for all population sizes. Game $\mathbf{G}_3$ also shows E[p] $> p_{\mathrm{Eq}_N}$; except for $N \in \{372, 756, 900\}$, $\overline{p}$ deviates with statistical significance. In game $\mathbf{G}_4$, we again find no statistically significant deviation.

5.2 Moran Replication

The four panels of Figure 4 give Markov-chain and agent-based simulation results for our example games under the modified Moran process without agent

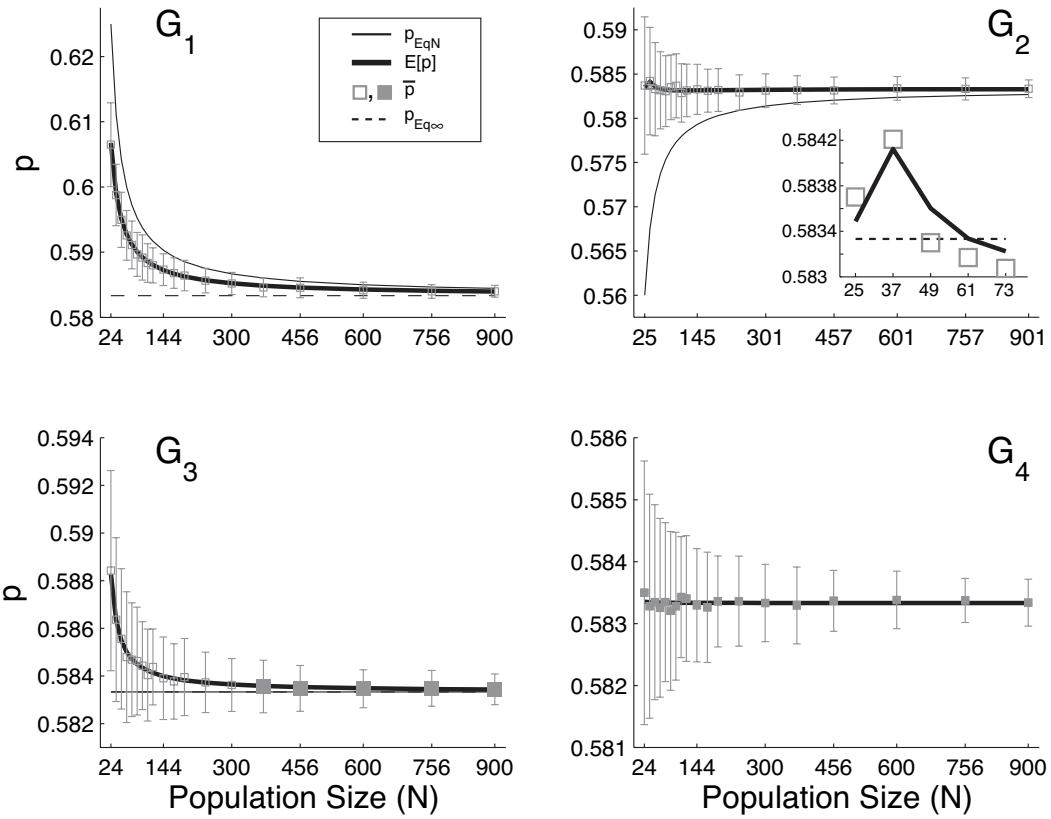

Fig. 2. Results for Wright-Fisher replicator dynamics without agent self-play. Thin solid curve indicates fitness-equilibrium p_{Eq_N} over different population sizes; bold solid curve indicates expected population-state E[p] (the mean of the pre-absorption distribution obtained with Markov model); boxes indicate mean population-state $\overline{p}$ observed from agent-based computer simulation, with standard deviation; empty boxes indicate statistically significant deviation away from p_{Eq_N}; dashed line indicates infinite-population fitness-equilibrium p_{Eq_∞}.

self-play. Qualitatively, the results are similar to the Wright-Fisher process (see Figure 2), except that the deviation away from fitness-equilibrium is greater under the Moran process (particularly in games $\mathbf{G}_1$ and $\mathbf{G}_2$). $\mathbf{G}_4$ again stands out as the only game that lacks statistically significant deviation between $\overline{p}$ and p_{Eq_N} over all population sizes. Figure 5 gives results under the Moran process with self-play allowed. These data fit the established pattern. The non-monotonic behavior of E[p] in game $\mathbf{G}_2$ is more pronounced.

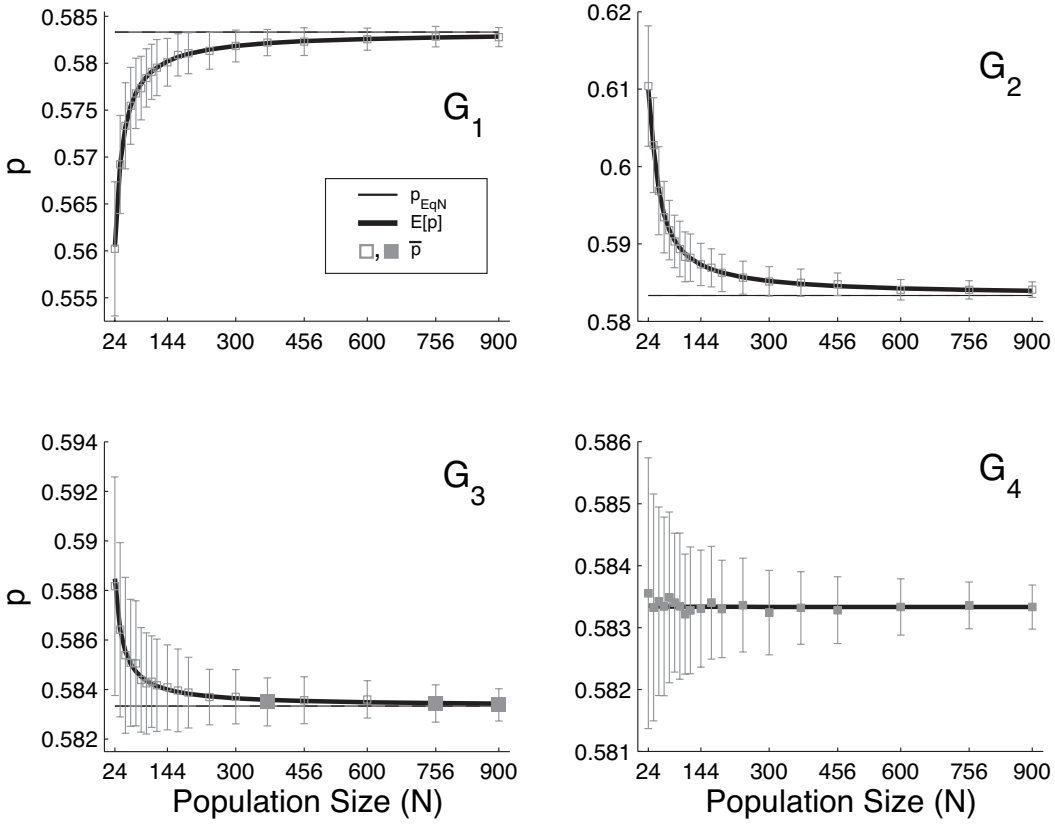

Fig. 3. Results for Wright-Fisher replicator dynamics with agent self-play. Thin solid line indicates fitness-equilibrium $p_{\mathrm{Eq}_N} = p_{\mathrm{Eq}_\infty} = 7/12$; bold solid curve indicates expected population-state $\mathrm{E}[p]$ (the mean of the pre-absorption distribution obtained with Markov model); boxes indicate mean population-state $\overline{p}$ observed from agent-based computer simulation, with standard deviation; empty boxes indicate statistically significant deviation away from p_{Eq_N}.

The increased divergence from p_{Eq_N}, compared with the Wright-Fisher process, is consistent with the smaller effective population size of the Moran process (Ewens, 2004; Wakeley and Takahashi, 2004). A cursory review of our data suggests that, with respect to divergence from fitness equilibrium, a population of size N under the Moran process acts similarly to a population of size $3N/5$ under Wright-Fisher.

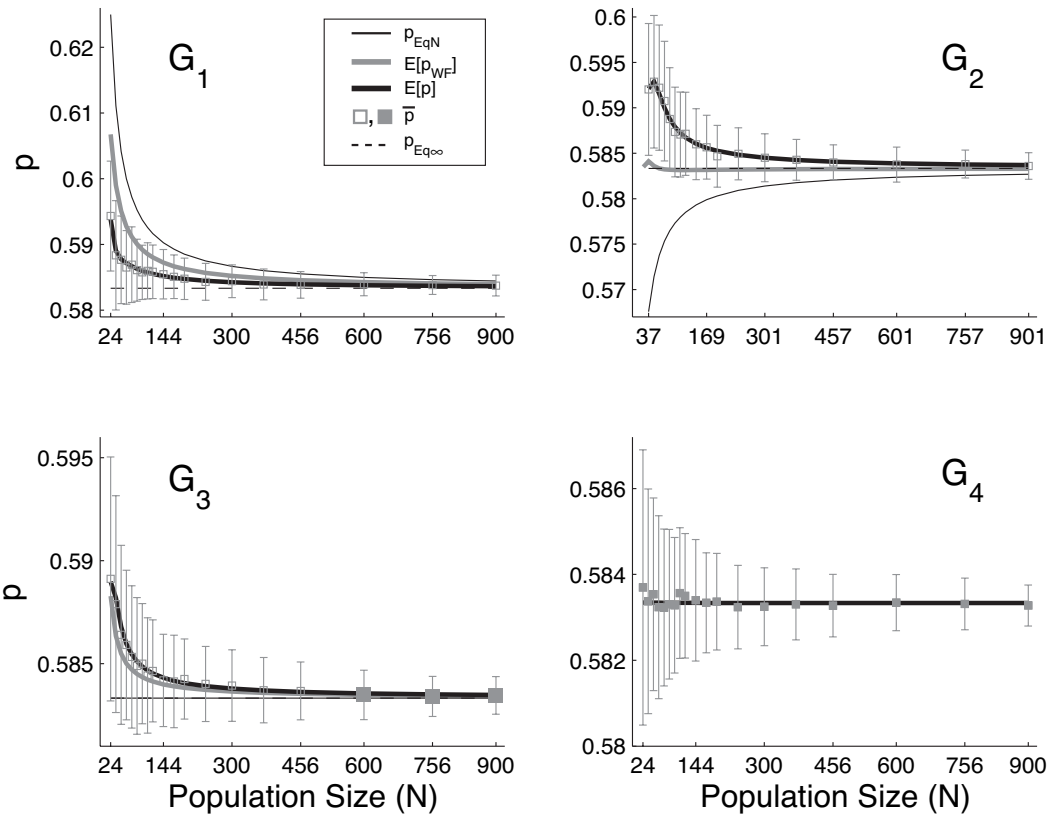

Fig. 4. Results for Moran replicator dynamics without agent self-play. Thin black curve indicates fitness-equilibrium p_{Eq_N} over different population sizes; bold black curve indicates expected population-state E[p] (the mean of the pre-absorption distribution obtained with Markov model); boxes indicate mean population-state $\overline{p}$ observed from agent-based computer simulation, with standard deviation; empty boxes indicate statistically significant deviation away from p_{Eq_N}; bold grey curve indicates expected population-state E[p_{WF}] (calculated by Markov model) under Wright-Fisher dynamics without agent self-play (indicated in Figure 2 as E[p]); dashed line indicates infinite-population fitness-equilibrium p_{Eq_∞}.

5.3 Summary of Empirical Results

Table 1 summarizes our empirical results for games $\mathbf{G}_1$ through $\mathbf{G}_4$. All four games share the same polymorphic fitness-equilibrium when played by an infinite population or by a finite population with self-play. When played by a

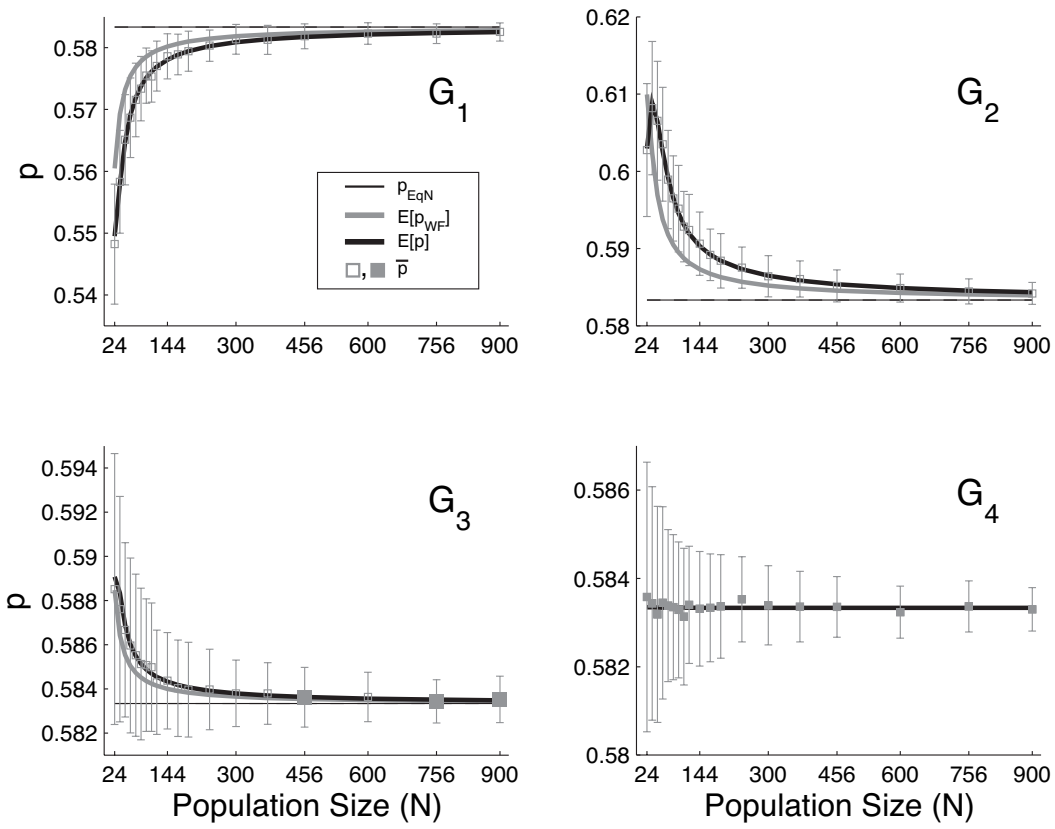

Fig. 5. Results for Moran replicator dynamics with agent self-play. Thin black line indicates fitness-equilibrium $p_{\mathrm{Eq}_N} = p_{\mathrm{Eq}_\infty} = 7/12$; bold black curve indicates expected population-state $\mathrm{E}[p]$ (the mean of the pre-absorption distribution obtained with Markov model); boxes indicate mean population-state $\overline{p}$ observed from agent-based computer simulation, with standard deviation; empty boxes indicate statistically significant deviation away from p_{Eq_N}; bold grey curve indicates expected population-state $\mathrm{E}[p_{\mathrm{WF}}]$ (calculated by Markov model) under Wright-Fisher dynamics with agent self-play (indicated in Figure 3 as $\mathrm{E}[p]$).

finite population without self-play, the polymorphic fitness-equilibrium p_{Eq_N} depends upon population size for games $\mathbf{G}_1$ and $\mathbf{G}_2$, but not $\mathbf{G}_3$ and $\mathbf{G}_4$ (see Figure 1).

Nevertheless, the mean population-states $\overline{p}$ observed from agent-based simulations, as well as the expected population-states $\mathrm{E}[p]$ determined from Markov-

chain models, generally do not correspond with the polymorphic fitness-equilibria of our games. We find that $\overline{p}$ diverges with statistical significance from p_{Eq_N} for game $\mathbf{G}_1$ (where $\overline{p} < p_{\mathrm{Eq}_N}$) and games $\mathbf{G}_2$ and $\mathbf{G}_3$ (where $\overline{p} > p_{\mathrm{Eq}_N}$), but not for game $\mathbf{G}_4$.

Figures 4 and 5 show that divergence from p_{Eq_N} is greater under the Moran process than under Wright-Fisher. We also observe, though our graphs do not make this conspicuous, that divergence is greater with self-play than without; the impact of self-play on divergence appears to be less than that caused by the Moran process.

Thus, our finite-population experiments elicit a variety of relationships between $\overline{p}$ and p_{Eq_N}, depending upon the game, the replication process, and whether self-play is used. Nevertheless, we will show below that this variety is superficial in this sense: the different behaviors we observe are all manifestations of a single underlying dynamic whereby our finite populations evolve to equilibrate selection pressure, not fitness.

6 Selection-Pressure Asymmetry

We now present our hypothesis that asymmetry in selection pressure can explain the empirical results detailed above. The intuitions we develop here will be formalized in Section 7.

Let $\Delta(p)$ be the expected change in population state when the replicator function $f(p)$ acts upon a population in state p, that is,

Table 1

Summary of simulation results. Top row gives infinite-population fitness-equilibrium for each game; second row gives finite-population fitness-equilibrium under self-play; third row gives finite-population fitness-equilibrium under no self-play ($\searrow$ and $\nearrow$ denote asymptotic approach from above and below, respectively—see Figure 1); bottom row indicates how mean and expected population-states relate to fitness-equilibrium over range of population sizes.

	$\mathbf{G}_1$	$\mathbf{G}_2$	$\mathbf{G}_3$	$\mathbf{G}_4$
p_{Eq_∞}	7/12	7/12	7/12	7/12
p_{Eq_N} (self-play)	7/12	7/12	7/12	7/12
p_{Eq_N} (no self-play)	$p_{\mathrm{Eq}_N} \searrow p_{\mathrm{Eq}_\infty}$ as $N \to \infty$	$p_{\mathrm{Eq}_N} \nearrow p_{\mathrm{Eq}_\infty}$ as $N \to \infty$	7/12	7/12
$\overline{p}$ and $\mathrm{E}[p]$	$< p_{\mathrm{Eq}_N}$	$> p_{\mathrm{Eq}_N}$	$> p_{\mathrm{Eq}_N}$	$\approx p_{\mathrm{Eq}_N}$

$$\Delta(p) = f(p) - p. \tag{12}$$

When selection favors X-strategists, the expected change in population state is positive and $\Delta(p) > 0$; similarly, when selection favors Y-strategists, $\Delta(p) < 0$. Thus, the sign of $\Delta(p)$ indicates the direction in which selection points when the population is in state p; the magnitude of $\Delta(p)$ indicates the strength of selection pressure being applied to the population.

Figure 6 plots the delta-function $\Delta(\cdot)$ for each of our example games, using the Wright-Fisher replicator dynamic operating on an infinite population (6). In each game, we see $\Delta(p_{\mathrm{Eq}_\infty} = 7/12) = 0$, $\Delta(0 < p < p_{\mathrm{Eq}_\infty}) > 0$, and $\Delta(p_{\mathrm{Eq}_\infty} < p < 1) < 0$; that is, selection pressure always points towards fitness-equilibrium.

Nevertheless, we find the magnitude of $\Delta(\cdot)$ (and thus selection pressure) to be very asymmetric about p_{Eq_∞}. In game $\mathbf{G}_1$, for example, $\Delta(\cdot)$ tends to have smaller magnitude below p_{Eq_∞} than it does above; thus, the rate at which the population approaches p_{Eq_∞} is slower (on average) from below equilibrium than from above. The opposite is true for $\mathbf{G}_2$.

The trajectory of a finite population is constantly buffeted by sampling error, which prevents the population from fixing onto p_{Eq_N}. As the population fluctuates from one side of p_{Eq_N} to the other, it will tend to spend more time on the side with weaker selection pressure. Consequently, the imbalance in selection pressure pushes the mean population-state away from fitness equilibrium.

The above intuition is consistent with our results. Our data fall below p_{Eq_N} in game $\mathbf{G}_1$, and above in $\mathbf{G}_2$ and $\mathbf{G}_3$. In game $\mathbf{G}_4$, our data do not deviate from equilibrium, and we see that $\Delta(\cdot)$ is very symmetric in this game. Thus, the game essentially determines the degree and shape of selection-pressure asymmetry that exists. Nevertheless, the amount of asymmetry to which the population is actually exposed is determined by the population's size. The Wright-Fisher process produces a binomial distribution in the interval $[0, 1]$ with expected value $f(p)$ and variance $p(1-p)/N$; variance grows as population size N shrinks. Larger variance increases exposure to the asymmetry in $\Delta(\cdot)$.

Equations (13) and (14) give the delta function for the Wright-Fisher replicator under self-play (alternatively, an infinite population) and no self-play, respectively:

$$\Delta(p) = \frac{p^2(a-b) + p(b + w_0 - \alpha)}{\alpha} \tag{13}$$

$$\Delta(p) = \frac{N\Big(p^2(a-b) + p(b + w_0 - \alpha)\Big) + (p - p^2)(d-a)}{N\alpha + p(d-a) - d - w_0} \tag{14}$$

$$\alpha = p^2(a - b - c + d) + p(b + c - 2d) + d + w_0$$

Note that (14) approaches (13) as $N \to \infty$. Simple algebra shows that the delta equations for the Moran replicator, under self-play and no self-play, are related to the corresponding Wright-Fisher equations by a constant factor that is equal to population size:

$$\Delta_{\text{Wright-Fisher}}(p) = N \cdot \Delta_{\text{Moran}}(p).$$

Selection-pressure asymmetry is therefore identical for the two replicator processes; the only difference is the absolute scale of the delta-values.

Nevertheless, the Moran process operates differently than the Wright-Fisher process. In Wright-Fisher, all N offspring are generated using the same random variable (with $f(p)$ as the expected value); in the Moran process, the N consecutive offspring are (almost certainly) generated using several different random variables (with different expected values). Thus, the effects of sampling error are compounded under the Moran process, which relates to the smaller effective population size and increased divergence from p_{Eq_N} of the Moran process.

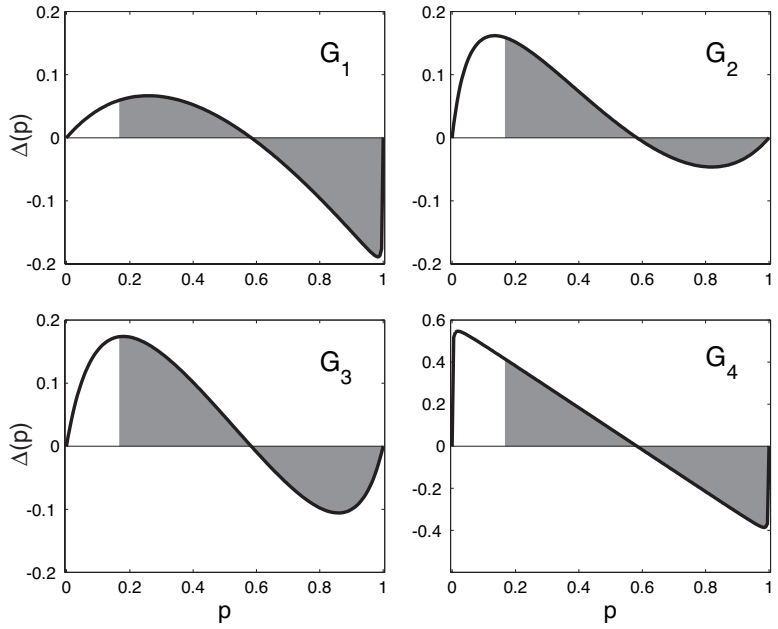

Fig. 6. $\Delta(\cdot)$ for games $\mathbf{G}_1$ through $\mathbf{G}_4$ when played by an infinite population. The x-axis indicates the population state p; the y-axis indicates the expected change $\Delta(p)$ in population state after one iteration of replicator equation (6). Shaded regions indicate integrals of $\Delta(\cdot)$ over intervals $[p_{\mathrm{Eq}_\infty} - \epsilon, p_{\mathrm{Eq}_\infty}]$ and $[p_{\mathrm{Eq}_\infty}, p_{\mathrm{Eq}_\infty} + \epsilon]$ for $p_{\mathrm{Eq}_\infty} = 7/12$ and $\epsilon = 5/12$. Note that for games $\mathbf{G}_1$, $\mathbf{G}_2$, and $\mathbf{G}_3$ the shaded regions are asymmetric, which indicates that selection pressure is asymmetric. For example, in $\mathbf{G}_1$, selection pressure for $p > p_{\mathrm{Eq}_\infty}$ (where selection favors Y-strategists) tends on average to be stronger than selection pressure for $p < p_{\mathrm{Eq}_\infty}$ (where selection favors X-strategists).

7 Selection Equilibrium

In this section, we make our intuitions about selection-pressure asymmetry concrete. Figure 7 shows $\Delta(\cdot)$ for game $\mathbf{G}_1$ under the Wright-Fisher process without self-play (7) acting on a population of size $N = 24$. Here, fitness-equilibrium is achieved at $p_{\mathrm{Eq}_N} = 0.625$ (15 X-strategists and 9 Y-strategists). The dashed curve indicates the binomial distribution produced by the replicator process when the population state is p_{Eq_N}; the expected value of this

distribution is also p_{Eq_N} and is indicated by the dashed vertical line. Thus, $\Delta(p_{\mathrm{Eq}_N}) = 0$. We use the Markov-chain model (see Appendix B) to calculate the expected number of visits each transient state receives before absorption occurs at time $t = T_{\mathrm{A}} + 1$. This pre-absorption distribution over the transient states is indicated in Figure 7 by the solid curve; the expected value of this distribution is $\mathrm{E}[p] = 0.607$—roughly two percent below p_{Eq_N}.

If the selection-pressure asymmetry indicated by $\Delta(\cdot)$ causes the expected population-state $\mathrm{E}[p]$ to diverge from the p_{Eq_N}, then we may expect $\mathrm{E}[p]$ to represent some kind of equilibrium with respect to this asymmetry. To test this hypothesis, we measure the mean selective pressures that are applied to a population from above and below p_{Eq_N} and calculate their ratio. Let $\lambda(t)$ be the quotient of selection pressures applied to the population by time t:

$$\lambda(t) = \frac{\sum_{p \geq p_{\mathrm{Eq}_N}} |\Delta(p)| V(p|t)}{\sum_{p \leq p_{\mathrm{Eq}_N}} \Delta(p) V(p|t)}. \tag{15}$$

The calculation of $\lambda(t)$ involves two weighted integrals. The numerator of (15) represents the expected cumulative selection pressure applied to the population by time t from above the polymorphic fitness-equilibrium. Specifically, for each population-state p at and above p_{Eq_N} (that a population of size N can visit), we take the absolute value of $\Delta(p)$, which represents the strength of selection, and weight it by the expected number of visits $V(p|t)$ to that state by time t; we then sum these products to form the numerator of the quotient. We apply a similar process to population states at and below p_{Eq_N} to form the denominator. Figure 7 illustrates the quantities involved in the calculation of $\lambda(t = T_{\mathrm{A}})$.

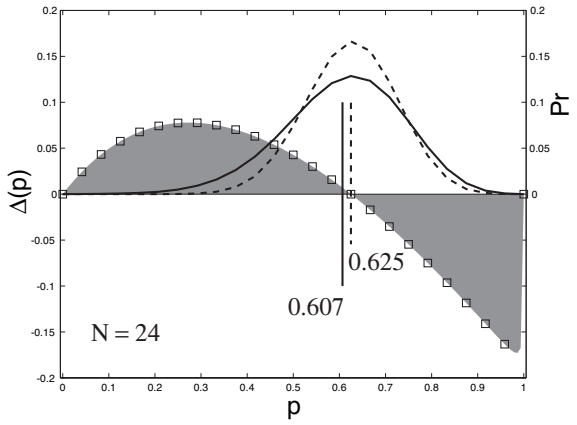

Fig. 7. Superimposition of $\Delta(\cdot)$ and probability distributions generated by replicator dynamics for game $\mathbf{G}_1$ with population size $N = 24$ under the Wright-Fisher replicator dynamic without agent self-play (7). The x-axis represents the population state p; the left y-axis represents $\Delta(p)$; the right y-axis represents probability. The shaded area indicates $\Delta(p)$ over the interval $p \in [0, 1]$, while the boxes indicate $\Delta(p)$ for those values of p that a population of size $N = 24$ can have, i.e., $p \in \{0, 1/24, 2/24, \ldots, 1\}$. Dashed curve indicates binomial distribution created by replicator process when population is at $p_{\mathrm{Eq}_N} = 0.625$; dashed vertical line indicates the mean of this distribution, which is also p_{Eq_N}. Solid curve indicates proportion of time spent in each population state before absorption occurs, as calculated by the Markov-chain model; solid vertical line indicates the mean of this pre-absorption distribution. We see that the average population state of the pre-absorption transient is below the polymorphic fitness-equilibrium; thus, the finite population does not equilibrate fitness in this game. If we instead weight $\Delta(\cdot)$ by the pre-absorption distribution (solid curve) and integrate below and above p_{Eq_N}, we find that the pre-absorption distribution equilibrates the asymmetries in $\Delta(\cdot)$; that is, the finite population does equilibrate selection pressure.

If the pre-absorption distribution $V(\cdot|T_{\mathrm{A}})$ equilibrates the asymmetries of $\Delta(\cdot)$, then we expect $\lambda(T_{\mathrm{A}}) = 1.0$. Figure 8 (top) shows the time evolution of the Markov chain for Wright-Fisher dynamics without self-play (7) on game $\mathbf{G}_1$

for several population sizes. As time t moves forward, $\lambda(t)$ asymptotically converges onto 1.0. The expected population state moves from $\mathrm{E}[p|1] = p_{\mathrm{Eq}_N}$ (which equilibrates fitness) towards $\mathrm{E}[p|T_{\mathrm{A}}]$, which is the mean of a distribution over population states $V(\cdot|T_{\mathrm{A}})$ that equilibrates $\Delta(\cdot)$ and selection pressure.

Note that $\lambda(t)$ approaches 1.0 from above; given the binomial distribution produced when the population is at p_{Eq_N}, the weighted integral of $\Delta(p)$ for states $p \geq p_{\mathrm{Eq}_N}$ is greater than the weighted integral for states $p \leq p_{\mathrm{Eq}_N}$. That is, excess selection pressure exists from above. As population size increases, the process' variance $p(1-p)/N$ decreases; this exposes the system to less of the asymmetry in $\Delta(\cdot)$ and decreases the initial value of $\lambda(t)$.

Figure 8 (bottom) shows the behavior of $\lambda(t)$ over time for $\mathbf{G}_1$ under the Moran process without self-play (11). We again see $\lambda(t)$ converge to 1.0 as the population converges onto a distribution that equilibrates the asymmetries of $\Delta(\cdot)$. Unlike with the Wright-Fisher process, we find that initial values of $\lambda(t)$ are near 1.0 (e.g., 1.035, 1.0197, and 1.0089 for population sizes $N = 60$, 108, and 240, respectively). This effect is due to the fact that the probability mass spreads more slowly per iteration under the Moran process than under the Wright-Fisher process; if we instead extract values from the Moran process every N iterations, then we obtain plots very similar to those of the Wright-Fisher process.

Since Equations (5)–(11) are mean-field equations, they do not model activity away from the mean; therefore, they cannot capture the asymmetry in selection pressure that underlies our result. Statistically, the behavior of a finite population converges onto a distribution over population states that equili-

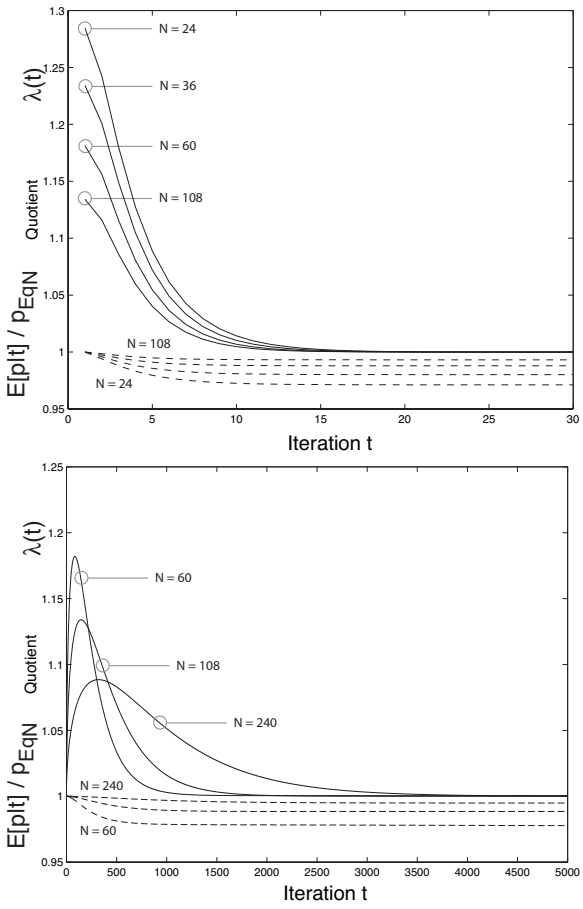

Fig. 8. Time-evolution of $\lambda(t)$ and $\mathrm{E}[p|t]/p_{\mathrm{Eq}_N}$ for Wright-Fisher (top) and Moran (bottom) replicator dynamics, calculated by iteration of reducible Markov-chain models. Finite populations of various sizes play game $\mathbf{G}_1$ without self-play. The expected population state $\mathrm{E}[p|t]$ diverges from p_{Eq_N}; thus, $V(\cdot|T_{\mathrm{A}})$ does not equilibrate fitness in this game. At the same time, $\lambda(t)$ converges onto 1.0; thus, the pre-absorption distribution $V(\cdot|T_{\mathrm{A}})$ over population states equilibrates the asymmetries in $\Delta(\cdot)$.

brates asymmetries in the selection pressures that act upon the population. To formalize our argument and generalize this result to other games, we focus on the Moran process, which has a simple Markov-chain model. By a small modification to the Markov chain to make it irreducible, we are able to obtain a general proof that $\lambda(\infty) = 1.0$ at the Markov chain's steady-state distribution

(see Appendix C). This proof applies to any game with the payoff structure $a < c$ and $b > d$.

Though the fitness-equilibrium state p_{Eq_N} does not accurately describe the expected population-state $\mathrm{E}[p]$, the fitness equilibrium is crucial to our understanding of the expected behavior. Recall that (15) uses p_{Eq_N} as the split-point in $\Delta(\cdot)$ to calculate $\lambda(t)$. Our proof shows that, for any game with the payoff structure $a < c$ and $b > d$, p_{Eq_N} is the unique split-point that always gives $\lambda(\infty) = 1.0$. Thus, p_{Eq_N} is key to understanding that $\mathrm{E}[p]$ is the mean of a distribution that equilibrates selection pressure.

8 Discussion

8.1 Main Findings

In Section 5, we demonstrate that different games can induce different behaviors in a finite population despite sharing the same infinite- or finite-population fitness-equilibrium. Three of the four games we study cause the mean population-state $\overline{p}$ to diverge with statistical significance from the finite-population fitness-equilibrium p_{Eq_N}; depending upon the game, the mean population-state is either too high or too low. With the fourth game, $\overline{p}$ does not diverge appreciably from p_{Eq_N}. These results indicate that p_{Eq_N} poorly predicts that actual mean behavior of a finite population.

In Section 6, we propose that the divergence of $\overline{p}$ from p_{Eq_N} results from asymmetries in selection pressure, and define the function $\Delta(p)$ to help visualize these asymmetries. We find that $\overline{p}$ falls below p_{Eq_N} whenever the mean

selection pressure applied to states $p > p_{\mathrm{Eq}_N}$ is stronger than that for states $p < p_{\mathrm{Eq}_N}$; similarly, $\overline{p}$ rises above p_{Eq_N} when selection pressure has the opposite asymmetry.

In Section 7, we define $\lambda(t)$, which quantifies the ratio of selection pressures, from states above and below p_{Eq_N}, that have been cumulatively applied to the population by time t. We demonstrate that $\lambda(t)$ converges onto 1.0, indicating that the population generates a distribution over possible states that equilibrates asymmetries in selection pressure. The mean of this distribution is $\mathrm{E}[p]$, which is generally not equal to p_{Eq_N}. We prove in Appendix C that, for any game with payoffs $a < c$ and $b > d$, a well-mixed finite population will equilibrate selection pressure, but not necessarily fitness. Thus, a finite population cannot fix onto p_{Eq_N} due to stochastic selection, and has an expected population state $\mathrm{E}[p] \neq p_{\mathrm{Eq}_N}$ due to asymmetry in selection pressure; nevertheless, our definition of $\lambda(t)$ shows how p_{Eq_N} remains relevant to understanding the population's actual behavior.

8.2 Alternative Initial Conditions

The initial conditions considered in this article place all probability mass on the polymorphic fitness-equilibrium state p_{Eq_N}. Nevertheless, our results generalize to any initial condition provided that the time T_{E} required for the population to equilibrate $\Delta(\cdot)$ is sufficiently less than the expected time before absorption T_{A}. The times T_{E} and T_{A} are both functions of the Markov chain and initial condition. Appendix D discusses the relationship between T_{E} and T_{A} in detail; it shows that very small populations without mutation may absorb before the pre-absorption distribution is able to equilibrate $\Delta(\cdot)$.

8.3 Alternative Replicator Processes and Mutation

A well-mixed population of size N yields a Markov-chain model with $N + 1$ states, s_0 through s_N, where state s_i represents a population with i X-strategists. The Markov chain used in our proof in Appendix C was derived from the Moran process defined in Section 3.5. Nevertheless, our proof applies to any replication process in which 1) some state s_i for $0 < i < N$ represents the polymorphic fitness-equilibrium p_{Eq_N} and 2) transitions other than self-loops occur only between states s_i and s_j where $|i - j| = 1$. Thus, the proof generalizes beyond the specific Moran process we use in this paper; there exists an infinity of replicator processes that yield the Markov-chain structure of our proof. Our proof shows that the two transitions that leave the fitness-equilibrium state p_{Eq_N} must have the same probability. Aside from this fact, the proof makes no assumptions about the transition probabilities of the Markov chain; we are free to assume that the transition probabilities subsume any arbitrary mutation process that fits the chain's structure. Thus, the addition of mutation will not change the result that $\lambda(t)$ converges to one.

8.4 Current Work

This article confines itself to one-dimensional systems. Nevertheless, we can easily construct a symmetric variable-sum game of $n > 2$ strategies that has a single polymorphic fitness-equilibrium attractor with all n strategies in support. Our preliminary work in higher dimensions shows that selection-pressure asymmetry remains ubiquitous and may even be more pronounced.

Another area of preliminary work concerns the delta function $\Delta(\cdot)$ generated

by the replicators we study. Specifically, when $\Delta(\cdot)$ is locally symmetric about the polymorphic fitness-equilibrium, then $\Delta(\cdot)$ must have an inflection point at the equilibrium, as well. Thus, a simple predictor of divergence from polymorphic fitness-equilibrium appears to be the location of the zero of the second derivative of $\Delta(\cdot)$; if the zero is not at the fitness-equilibrium state, then $\Delta(\cdot)$ is asymmetric and deviation will be observed. The relationship between the location of the inflection point and the degree of deviation remains to be closely studied.

Given a polymorphic fitness-equilibrium at p_{Eq_N}, what is the maximal deviation away from p_{Eq_N} that we can observe in each direction? Though most of the deviations we observe in this paper are statistically significant, they are small in magnitude and usually represent less than one individual. We are investigating methods to compute bounds.

Finally, this article only considers well-mixed populations. A natural extension of this work is to examine how spatially structured populations react to asymmetries in selection pressure.

9 Conclusion

This article considers simple variable-sum games that have polymorphic fitness-equilibrium attractors (under standard replicator dynamics) when played by an infinite population of pure-strategists. When these games are played by finite, fixed-size populations of pure-strategists, we show that the expected population state will likely diverge from the polymorphic fitness-equilibrium; data for Wright-Fisher and frequency-dependent Moran replication processes

are given. We then show that this deviation occurs when the selection pressures that surround the polymorphic fitness-equilibrium are asymmetric. Further, we prove that the expected population state obtained under these replicator dynamics is the mean of a distribution that equilibrates this selection-pressure asymmetry.

Acknowledgments

The authors thank Bruce Boghosian, David Fogel, David Haig, Lorens Imhof, Martin Nowak, Shivakumar Viswanathan, John Wakeley, Daniel Weinreich, L. Darrell Whitley, and the anonymous reviewers for their helpful feedback on this work. This work was supported in part by the U.S. Department of Energy (DOE) under Grant DE-FG02-01ER45901.

A Agent-Based Simulations

Let $\mathbf{G}$ be a symmetric 2x2 game with payoff structure $a < c$ and $b > d$ (see Section 3.1). We assume a background fitness $w_0 = 0$ and a population of fixed size N. Let $\mathbf{G}$ have a polymorphic fitness-equilibrium at p_{Eq_N}. If agent self-play is allowed, then we use (2) to calculate p_{Eq_N}; otherwise, we use (4). We require that $N \cdot p_{\mathrm{Eq}_N}$ be a whole number so that we can initialize the population to be precisely at fitness equilibrium.

Each trial of an experiment consists of T time-steps, or iterations of the simulation. We count the number of times each population state is entered during the trial. We begin each trial with a population of $N \cdot p_{\mathrm{Eq}_N}$ X-strategists

and $N \cdot (1 - p_{\mathrm{Eq}_N})$ Y-strategists. If the system absorbs before T time-steps pass, then we re-initialize the population and continue the trial, adding to our population-state counts. (Note that we do not include the initial condition at the beginning of a trial or at re-initialization in our counts.) Using our counts of population-state visits, we calculate the mean population state (expressed as the proportion of X-strategists in the population) and note the result. We then zero our counts and perform another trial of the same experiment. Each experiment consists of K trials. The results we report concern the mean and standard deviation of these K trials for each experiment. For experiments that use Wright-Fisher replication, we set $T = 2000$ and $K = 200$; experiments that use Moran replication have $T = 2000N$ and $K = 200$. These parameter values are large enough to produce accurate estimations of mean population state for the population sizes we examine; as shown in Figures 2–5, our estimations agree very well with the expected population states computed with our Markov-chain models.

B Reducible Markov-Chain Models

A population of N pure-strategists playing a symmetric 2x2 game can be modeled with $N + 1$ Markov states. Let s_i, for $i = 0 \ldots N$, be the state representing a population with i X-strategists. Since we lack mutation, states s_0 and s_N are absorbing and $s_{1 \ldots N-1}$ are transient. Let M be the $N+1$ by $N+1$ transition matrix. $M(i, j)$ is the transition probability from state s_i to s_j; each row of M sums to 1.0. The entries of M are calculated with Equations (B.1) and (B.2) for the Wright-Fisher and Moran processes, respectively. Equations (1) and (3) are used to calculate cumulative payoffs (w_{X} and w_{Y}) for self-play

and no self-play, respectively.

The transition matrix M that we obtain can be decomposed as follows:

$$M = \begin{bmatrix} 1 & \mathbf{0} & 0 \\ * & \mathbf{Q} & * \\ 0 & \mathbf{0} & 1 \end{bmatrix}$$

The top row vector [1**0**0] and bottom row vector [0**0**1] signify that the absorbing states s_0 and s_N, respectively, can only transition to themselves. The $*$ is used to denote column vectors that represent transition probabilities from the transient states $s_{1\ldots N-1}$ to the absorbing states. The submatrix $\mathbf{Q}$ represents the transition probabilities from transient states to transient states.

To calculate the expected number of visits each transient state receives before absorption occurs, we use the *fundamental matrix* of M, which is defined to be $F_M = [\mathbf{I} - \mathbf{Q}]^{-1}$, where $\mathbf{I}$ is the identity matrix and the exponent -1 refers to matrix inversion. $F_M(k, l)$ indicates the expected number of visits transient state s_{l+1} receives when the Markov chain is begun in transient state s_{k+1}, for $0 \leq k, l \leq N-2$. The sum of elements in row k of F_M thus gives us the expected time before absorption T_{A} when the Markov chain is begun in transient state s_{k+1}. We compute the fundamental matrix of M by the method of Sheskin (1997), which is resistant to round-off error and is numerically stable even for large transition matrices.

$$M(i,j) = \binom{N}{j} q^j (1-q)^{N-j} \qquad \text{for } 0 \leq i,j \leq N \tag{B.1}$$

$$M(i,j) = \begin{cases} q(1-i/N) & \text{for } 0 < i < N \text{ and } j = i+1 \\ (1-q)i/N & \text{for } 0 < i < N \text{ and } j = i-1 \\ 1 - q(1-i/N) - (1-q)i/N & \text{for } j = i \\ 0 & \text{otherwise} \end{cases} \tag{B.2}$$

$$q = \frac{i \cdot w_{\mathrm{X}}}{i \cdot w_{\mathrm{X}} + (N-i) \cdot w_{\mathrm{Y}}}$$

C Proof that Finite Populations Equilibrate Selection Pressure

Let $\mathcal{M}$ denote the reducible Markov-chain model for the frequency-dependent Moran process specified by Equation (B.2). Given a population of size N, the reducible Markov chain $\mathcal{M}$ has $N+1$ states. Let s_i be the state with i X-strategists. States s_0 and s_N are absorbing states. Let M be the transition matrix and $M(i,j)$ denote the transition probability from state s_i to s_j. Let $\mathcal{M}'$ denote the irreducible Markov chain that corresponds to $\mathcal{M}$, and let M' denote the transition matrix of $\mathcal{M}'$. Here we prove that $\lambda(t=\infty) = 1.0$ at the steady-state distribution of $\mathcal{M}'$.

We present two methods to modify $\mathcal{M}$ to obtain $\mathcal{M}'$. Our first method (which we use in Appendix D) creates $\mathcal{M}'$ by removing the absorbing states s_0 and s_N of $\mathcal{M}$, leaving a system of $N-1$ recurrent states $s_{1\ldots N-1}$. Since states

s_0 and s_N no longer exist, we obtain the self-loop probabilities $M'(1,1)$ and $M'(N-1,N-1)$ as follows:

$$M'(1,1) = M(1,1) + M(1,0),$$
$$M'(N-1,N-1) = M(N-1,N-1) + M(N-1,N).$$

Our second method creates $\mathcal{M}'$ by introducing mutation at the monomorphic states s_0 and s_N of $\mathcal{M}$; this makes the monomorphic states elastic, rather than absorbing. Each monomorphic state will transition to its neighboring polymorphic state with a non-zero probability:

$$M'(0,1) = \rho_1,$$
$$M'(0,0) = 1.0 - \rho_1,$$
$$M'(N,N-1) = \rho_2,$$
$$M'(N,N) = 1.0 - \rho_2,$$

where

$$0 < \rho_1, \rho_2 < 1.0.$$

Our proof is indifferent to the particulars of the transformation from $\mathcal{M}$ to $\mathcal{M}'$. With either method, we obtain an irreducible Markov-chain where each state has a self-loop, each interior state can transition to its immediate left and right neighbors, and the states at the ends of the chain can transition to their neighboring interior states. Without loss of generality, Figure C.1 illustrates $\mathcal{M}'$ with seven states. We arbitrarily select interior state s_3 to be

the polymorphic fitness-equilibrium state, which we denote s_E. Though $\mathcal{M}'$ has self-loops, they are not involved in the proof and so are omitted from the diagram.

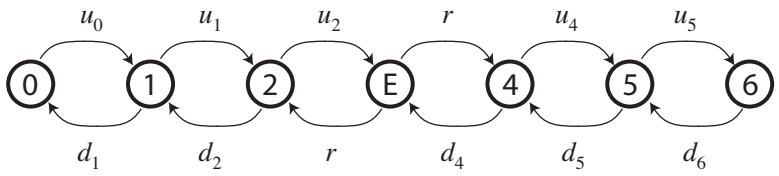

Fig. C.1. State diagram of irreducible Markov-chain $\mathcal{M}'$.

Remark 1 (Steady-State Distribution) *In the steady-state distribution of $\mathcal{M}'$, the probability of being in state s_i can be calculated as follows (adapted from Taylor and Karlin, 1998, p. 248). For each state s_i, other than s_E, we calculate the product of transitions from s_E to s_i and divide by the product of transitions from s_i to s_E* (C.3)*; we call this quantity $g(s_i)$. For example, $g(s_1) = (r \cdot d_2)/(u_1 \cdot u_2)$. We then calculate* $\Pr(s_i)$ *as indicated in* (C.1) *and* (C.2).

$$\Pr(s_i) = \frac{g(s_i)}{1 + \sum_{s_j \neq s_E} g(s_j)} \qquad \textit{for } s_i \neq s_E \qquad \text{(C.1)}$$

$$\Pr(s_E) = 1 - \sum_{s_i \neq s_E} \Pr(s_i) = \frac{1}{1 + \sum_{s_j \neq s_E} g(s_j)} \qquad \textit{for } s_i = s_E \qquad \text{(C.2)}$$

$$g(s_i) = \frac{\prod s_E \rightsquigarrow s_i}{\prod s_i \rightsquigarrow s_E} \qquad \textit{for } s_i \neq s_E \qquad \text{(C.3)}$$

Lemma 2 *$M'(E, E+1) = M'(E, E-1)$; that is, the transition probabilities leaving s_E must be equal (both are labeled r in Figure C.1).*

PROOF. At state s_E we have $w_{\mathrm{X}} = w_{\mathrm{Y}}$, by definition. Thus, the probability of selecting an X-strategist to parent the offspring is simply the proportion

p_{Eq_N} with which it appears in the population; the probability of selecting a Y-strategist is $1-p_{\mathrm{Eq}_N}$. The probability of picking an X-strategist for replacement is also p_{Eq_N}; the probability of replacing a Y-strategist is $1-p_{\mathrm{Eq}_N}$. Thus, the probability of increasing the number of X-strategists is $p_{\mathrm{Eq}_N}(1-p_{\mathrm{Eq}_N})$; this is also the probability of decreasing the number of X-strategists. □

Theorem 3 *$\lambda(t=\infty)=1.0$ at the steady-state distribution of $\mathcal{M}'$.*

PROOF. To calculate $\lambda(t=\infty)$, we must know for each state s_i the value of $\Delta(s_i)$ and the probability $\Pr(s_i)$ of being in state s_i at the steady-state distribution of the Markov chain. The calculation of $\Pr(s_i)$ is shown above. Equations (C.4)–(C.6) show that $\Delta(s_i)$ is simply the difference between the probabilities of the two transitions that exit from state s_i, scaled by the population size N; we may safely discard this constant N.

$$f(p) = \Pr(X^0)p + \Pr(X^+)(p+\frac{1}{N}) + \Pr(X^-)(p-\frac{1}{N}) \tag{C.4}$$

$$\begin{aligned}\Delta(p) &= f(p) - p \\ &= \Pr(X^0)p + \Pr(X^+)(p+\frac{1}{N}) + \Pr(X^-)(p-\frac{1}{N}) - p \\ &= \Big(\Pr(X^+) - \Pr(X^-)\Big)/N \end{aligned} \tag{C.5}$$

$$\Delta(s_i) = \Big(M'(i,i+1) - M'(i,i-1)\Big)/N \tag{C.6}$$

When we calculate the numerator of $\lambda(\infty)$ (Equation (C.7) illustrates this calculation for the chain in Figure C.1), we find a telescoping sum; all terms are cancelled except for r. We find another telescoping sum when we calculate

the denominator (C.8); again, only r remains. This leaves $\lambda(\infty) = r/r = 1.0$. Because of the telescoping structure, we can arbitrarily extend the Markov chain in either direction without affecting the value of $\lambda(\infty)$. $\square$

$$
\begin{aligned}
|\Delta(s_{\mathrm{E}})|\Pr(s_{\mathrm{E}}) &= 0 \\
|\Delta(s_4)|\Pr(s_4) &= (d_4 - u_4)\frac{r}{d_4} \\
|\Delta(s_5)|\Pr(s_5) &= (d_5 - u_5)\frac{ru_4}{d_4 d_5} \\
|\Delta(s_6)|\Pr(s_6) &= (d_6 - 0)\frac{ru_4u_5}{d_4 d_5 d_6} \\
\sum_{i=3}^{6} |\Delta(s_i)|\Pr(s_i) &= r
\end{aligned}
\tag{C.7}
$$

$$
\begin{aligned}
\Delta(s_0)\Pr(s_0) &= (u_0 - 0)\frac{d_1 d_2 r}{u_0 u_1 u_2} \\
\Delta(s_1)\Pr(s_1) &= (u_1 - d_1)\frac{d_2 r}{u_1 u_2} \\
\Delta(s_2)\Pr(s_2) &= (u_2 - d_2)\frac{r}{u_2} \\
\Delta(s_{\mathrm{E}})\Pr(s_{\mathrm{E}}) &= 0 \\
\sum_{i=0}^{3} \Delta(s_i)\Pr(s_i) &= r
\end{aligned}
\tag{C.8}
$$

D Very Small Populations Without Mutation

Our proof in Appendix C shows why $\lambda(\infty) = 1.0$ for a well-mixed population under the Moran process with mutation. Empirical results presented in Section 7 indicate that $\lambda(t)$ also approaches 1.0 under Wright-Fisher and Moran processes that do not include mutation. Nevertheless, for very small popula-

tions without mutation, we observe that $\lambda(t)$ does not converge onto 1.0; for example, this occurs under the Moran process for game $\mathbf{G}_2$ with $N = 25$. To understand why this is so, we compare the Markov-chain models of the Moran process with and without mutation.

Let $\mathcal{M}$ denote the reducible Markov-chain model of the Moran process without mutation, and let $\mathcal{M}'$ denote the corresponding irreducible version of the Markov-chain used in our proof (see Appendix C for details on how we convert $\mathcal{M}$ to $\mathcal{M}'$). Two distinct dynamics are operating between these two Markov chains—each with their own speed—and they are in a race. First, there is the absorption time T_{A} of the reducible Markov-chain $\mathcal{M}$; this is the expected number of time-steps $\mathcal{M}$ will run before it enters one of the two absorbing population-states. Second, there is the "equilibration time" T_{E} of the irreducible Markov-chain $\mathcal{M}'$; this is the number of time-steps $\mathcal{M}'$ must run for $\lambda(t)$ to get within some ϵ of 1.0. (Our proof shows that $\lambda(\infty) = 1.0$ at the steady-state distribution of $\mathcal{M}'$.) If T_{E} is not sufficiently less than T_{A}, then the pre-absorption transient of $\mathcal{M}$ will not last long enough to equilibrate asymmetries in selection pressure; hence, $\lambda(T_{\mathrm{A}})$ does not reach 1.0.

Though the issue of very small populations applies generally to the class of games we study in this article, the specific population size beyond which $|1.0 - \lambda(T_{\mathrm{A}})| < \epsilon$ depends upon the game. For example, Figure D.1 compares the behaviors of $\mathcal{M}$ and $\mathcal{M}'$ on game $\mathbf{G}_2$ under the Moran process without self play. For the smaller population sizes, we clearly see that the pre-absorption transient produced by $\mathcal{M}$ has an expected population state $\mathrm{E}_{\mathcal{M}}[p|t = T_{\mathrm{A}}]$ that diverges from the expected population-state $\mathrm{E}_{\mathcal{M}'}[p|t = \infty]$ produced by the steady-state distribution of $\mathcal{M}'$; we also see that $\lambda(T_{\mathrm{A}})$ (calculated from the transient of $\mathcal{M}$) is not very near 1.0. As population size grows, the pre-

absorption transient of $\mathcal{M}$ lengthens and converges onto a distribution that equilibrates selection pressure, and $\lambda(T_{\mathrm{A}})$ goes to 1.0.

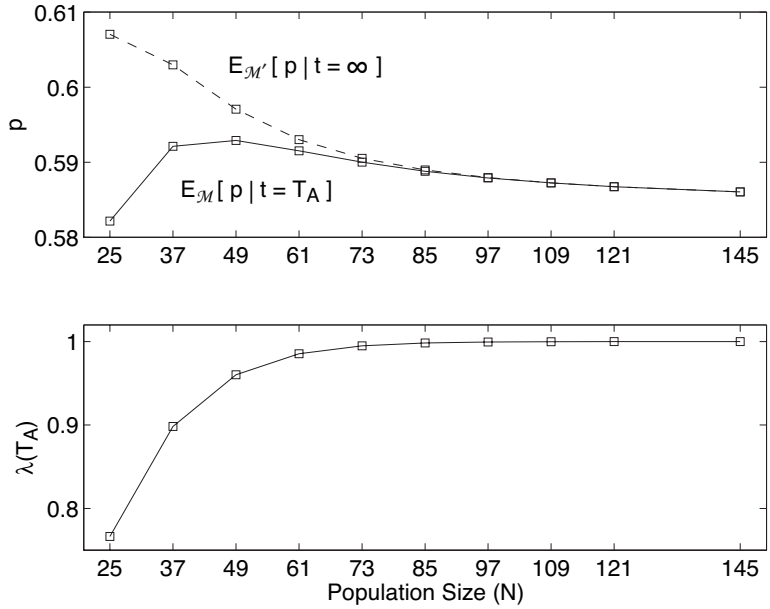

Fig. D.1. Behavior of Moran process without self-play in game $\mathbf{G}_2$ over different population sizes. Top: Expected population state $\mathrm{E}_{\mathcal{M}}[p|t = T_{\mathrm{A}}]$ produced by reducible Markov-chain $\mathcal{M}$ over pre-absorption transient compared with expected population-state $\mathrm{E}_{\mathcal{M}'}[p|t = \infty]$ produced by irreducible Markov-chain $\mathcal{M}'$ at steady-state. Bottom: $\lambda(T_{\mathrm{A}})$ calculated from pre-absorption transient of $\mathcal{M}$.

Figure D.2 illustrates the race between absorption and equilibration of $\Delta(\cdot)$ more directly. The top graph plots, for game $\mathbf{G}_2$ over different population sizes, the probability of having absorbed by iteration t (solid curves) and the expected time before absorption T_{A} (dashed vertical lines). (Note that the probability of absorption at $t = T_{\mathrm{A}}$ is approximately 0.632 for all four population sizes; we obtain this value with games $\mathbf{G}_1$, $\mathbf{G}_3$, and $\mathbf{G}_4$ over different population sizes, as well. Indeed, testing over a variety of games and population sizes, we find the probability of having absorbed by iteration $t = T_{\mathrm{A}}$ to be fairly consistent—from approximately 0.631 to 0.633—assuming we begin the population at fitness equilibrium.) The bottom graph plots for $\mathbf{G}_2$ the

time evolution of $\lambda(t)$ for the different population sizes, as calculated from the irreducible Markov-chain $\mathcal{M}'$. For population size $N = 25$, $\lambda(t)$ first exceeds 0.9997 at time-step $t = 843$, where $\Pr(\text{absorption}|t) \approx 0.327$. In contrast, for the larger population of $N = 61$, $\lambda(t)$ first exceeds 0.9997 at $t = 2026$, where $\Pr(\text{absorption}|t) \approx 0.0125$. Thus, larger populations provide $\mathcal{M}$ longer-lived transients, which in turn allow the system to better equilibrate the selection-pressure asymmetries of $\Delta(\cdot)$.

For sufficiently small populations, $\mathcal{M}$ is influenced more by the process of absorption than by the process of equilibrating asymmetries in selection pressure. Nevertheless, Figure D.2 shows that the relative strengths of these influences rapidly invert as population size increases: absorption time grows exponentially, whereas (dividing iterations by population size) T_{E} actually decreases. The decrease in T_{E} is consistent with the notion that larger populations expose the system to less of the extant asymmetry in $\Delta(\cdot)$. Thus, a sufficiently large population can spend a very long time in a quasi steady-state that equilibrates selection pressure before the system's probability mass shifts towards absorption.

E Schaffer's ESS Equation

An interesting comparison can be made between Equation (4) and the equation given by Schaffer (1988, Equation 16) to find an evolutionarily stable strategy (ESS) in a finite population, which we reproduce here:

$$s^{\mathrm{ESS}} = \frac{\frac{N-1}{N-2}b - d - \frac{1}{N-2}c}{b - d + c - a}.$$

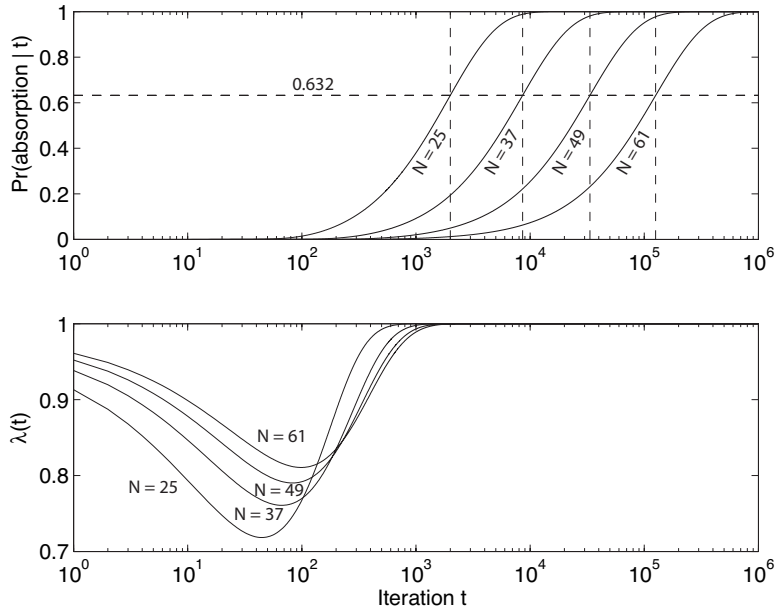

Fig. D.2. Race between absorption in $\mathcal{M}$ and equilibration of $\Delta(\cdot)$ in $\mathcal{M}'$. The population is playing game $\mathbf{G}_2$ under the frequency-dependent Moran process without self play for population sizes $N \in \{25, 37, 49, 61\}$. Top: Data obtained from the reducible Markov-chain $\mathcal{M}$. Solid curves indicate probability of having absorbed by iteration t; dashed vertical lines indicate expected time before absorption T_{A}. Dashed horizontal line indicates probability of having absorbed by time $t = T_{\mathrm{A}}$; this value is approximately 0.632 for all four population sizes. Bottom: Time evolution of $\lambda(t)$ for each population size; these data are obtained from the irreducible Markov-chain $\mathcal{M}'$.

The two equations are different, and indeed concern quite different situations. Schaffer (1988) assumes that agents are mixed-strategists, and he is interested to calculate an ESS strategy that resists invasion given a population of $N-1$ ESS agents and one arbitrary mutant agent; the ESS strategy obtained from his equation guarantees that the mutant agent (playing any mixture over the pure strategies in support of the ESS mixed-strategy) will obtain the same cumulative score (after complete mixing) as an ESS agent. In contrast, we assume a finite population of pure-strategist agents, and Equation (4) is used

to calculate the polymorphism proportions that allow pure strategists (playing either strategy X or Y) to obtain the same cumulative score.

Setting aside the fact that Equation (4) and Schaffer (1988, Equation 16) make different assumptions, we may still contrast their behaviors given identical inputs. Given the same game payoffs and population size, Equation (4) and Schaffer (1988, Equation 16) generally return different quantities. These equations will return the same quantity if and only if we meet the constraint that $N(b-c) = (N-2)(d-a)$. After some algebraic manipulation, we find that a combination of payoffs and population size will meet this constraint if and only if Schaffer (1988, Equation 16), or equivalently Equation (4), returns the value 0.5. Under this special circumstance, we have a situation where the mixed-strategy proportions of the ESS coincide with the fitness-equilibrium proportions between pure-strategists in a polymorphic population.

References

Benaïm, M., Schreiber, S. J., Tarrès, P., 2004. Generalized urn models of evolutionary processes. The Annals of Applied Probability 14 (3), 1455–1478.

Bergstrom, C. T., Godfrey-Smith, P., 1998. On the evolution of behavioral heterogeneity in individuals and populations. Biology and Philosophy 13, 205–231.

Ewens, W. J., 2004. Mathematical Population Genetics, 2nd Edition. Springer.

Ficici, S. G., Melnik, O., Pollack, J. B., 2000. A game-theoretic investigation of selection methods used in evolutionary algorithms. In: Zalzala, A., et al. (Eds.), Proceedings of the 2000 Congress on Evolutionary Computation.

IEEE Press, pp. 880–887.

Ficici, S. G., Melnik, O., Pollack, J. B., 2005. A game-theoretic and dynamical-systems analysis of selection methods in coevolution. IEEE Transactions on Evolutionary Computation 9 (6), 580–602.

Ficici, S. G., Pollack, J. B., 2000. Effects of finite populations on evolutionary stable strategies. In: Whitley, L. D., et al. (Eds.), Proceedings of the 2000 Genetic and Evolutionary Computation Conference. Morgan-Kaufmann, pp. 927–934.

Fisher, R. A., 1930. The genetical theory of natural selection. Clarendon Press.

Fogel, D. B., Fogel, G. B., 1995. Evolutionary stable strategies are not always stable under evolutionary dynamics. In: Evolutionary Programming IV. pp. 565–577.

Fogel, D. B., Fogel, G. B., Andrews, P. C., 1997. On the instability of evolutionary stable states. BioSystems 44, 135–152.

Fogel, G. B., Andrews, P. C., Fogel, D. B., 1998. On the instability of evolutionary stable strategies in small populations. Ecological Modelling 109, 283–294.

Hofbauer, J., Sigmund, K., 1998. Evolutionary Games and Population Dynamics. Cambridge University Press.

Lessard, S., 2005. Long-term stability from fixation probabilities in finite populations: New perspectives for ESS theory. Theoretical Population Biology 68, 19–27.

Liekens, A. M. L., ten Eikelder, H. M. M., Hilbers, P. A. J., 2004. Predicting genetic drift in 2x2 games. In: Deb, K., Poli, R., Banzhaf, W., Beyer, H.-G., Burke, E. (Eds.), Proceedings of the 2004 Genetic and Evolutionary Computation Conference. Lecture Notes in Computer Science 3102. Springer, pp. 549–560.

Maynard Smith, J., 1982. Evolution and the Theory of Games. Cambridge University Press.

Maynard Smith, J., 1988. Can a mixed strategy be stable in a finite population? Journal of Theoretical Biology 130, 247–251.

Moran, P., 1958. Random processes in genetics. Proceedings of the Cambridge Philosophical Society 54, 60–71.

Nowak, M. A., Sasaki, A., Taylor, C., Fudenberg, D., April 2004. Emergence of cooperation and evolutionary stability in finite populations. Nature 428, 646–650.

Orzack, S. H., Hines, W., 2005. The evolution of strategy variation: Will an ESS evolve? Evolution 59 (6), 1183–1193.

Riley, J. G., 1979. Evolutionary equilibrium strategies. Journal of Theoretical Biology 76, 109–123.

Schaffer, M. E., 1988. Evolutionarily stable strategies for a finite population and a variable contest size. Journal of Theoretical Biology 132, 469–478.

Schreiber, S. J., 2001. Urn models, replicator processes, and random genetic drift. SIAM Journal on Applied Mathematics 61, 2148–2167.

Sheskin, T. J., 1997. Computing the fundamental matrix for a nonirreducible Markov chain. International Journal of Mathematical Education in Science and Technology 28 (5), 661–675.

Taylor, C., Fudenberg, D., Sasaki, A., Nowak, M., 2004. Evolutionary game dynamics in finite populations. Bulletin of Mathematical Biology 66, 1621–1644.

Taylor, H., Karlin, S., 1998. An Introduction to Stochastic Modeling, 3rd Edition. Academic Press.

Taylor, P. D., Jonker, L., 1978. Evolutionarily stable strategies and game dynamics. Mathematical Biosciences 40, 145–156.

Vickery, W. L., 1987. How to cheat against a simple mixed strategy ESS. Journal of Theoretical Biology 127, 133–139.

Vickery, W. L., 1988. Reply to Maynard Smith. Journal of Theoretical Biology 132, 375–378.

Wakeley, J., Takahashi, T., 2004. The many-demes limit for selection and drift in a subdivided population. Theoretical Population Biology 66, 83–91.

Wright, S., 1931. Evolution in Mendelian populations. Genetics 16, 97–159.