



Research Article

Stochasticity, Selection, and the Evolution of Cooperation in a Two-Level Moran Model of the Snowdrift Game

Brian McLoone,^{1,2,3} Wai-Tong Louis Fan,^{1,4} Adam Pham,⁵
Rory Smead,⁶ and Laurence Loewe ^{1,2}

¹ Wisconsin Institute for Discovery, University of Wisconsin-Madison, 330 North Orchard Street, Madison, WI 53715, USA

² Laboratory of Genetics, University of Wisconsin-Madison, 425-G Henry Mall, Madison, WI 53706, USA

³ School of Philosophy, National Research University Higher School of Economics, Staraya Basmannaya st., No. 21/4, Moscow 105066, Russia

⁴ Department of Mathematics, University of Wisconsin-Madison, 408 Lincoln Drive, Madison, WI 53706, USA

⁵ Department of Philosophy, University of Wisconsin-Madison, 600 North Park Street, Madison, WI 53706, USA

⁶ Department of Philosophy and Religion, Northeastern University, 360 Huntington Avenue, 371 Holmes Hall, Boston, MA 02115, USA

Correspondence should be addressed to Laurence Loewe; loewe@wisc.edu

Received 1 September 2017; Accepted 28 November 2017; Published 13 February 2018

Academic Editor: Francisco J. S. Lozano

Copyright © 2018 Brian McLoone et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

The Snowdrift Game, also known as the Hawk-Dove Game, is a social dilemma in which an individual can participate (cooperate) or not (defect) in producing a public good. It is relevant to a number of collective action problems in biology. In a population of individuals playing this game, traditional evolutionary models, in which the dynamics are continuous and deterministic, predict a stable, interior equilibrium frequency of cooperators. Here, we examine how finite population size and multilevel selection affect the evolution of cooperation in this game using a two-level Moran process, which involves discrete, stochastic dynamics. Our analysis has two main results. First, we find that multilevel selection in this model can yield significantly higher levels of cooperation than one finds in traditional models. Second, we identify a threshold effect for the payoff matrix in the Snowdrift Game, such that below (above) a determinate cost-to-benefit ratio, cooperation will almost surely fix (go extinct) in the population. This second result calls into question the explanatory reach of traditional continuous models and suggests a possible alternative explanation for high levels of cooperative behavior in nature.

1. Introduction

Evolutionary game theory (Hofbauer and Sigmund 1998 [1]; Smith 1982 [2]; Smith and Price 1973 [3]) allows one to analyze the evolutionary dynamics of a population of individuals in which individual fitness is frequency dependent. A model in evolutionary game theory is based on a payoff matrix, which describes the payoff an individual will receive given its own behavior and the behavior of its partner(s). Payoff Matrix 1 represents a symmetric, two-player game. The entries in the matrix— α , β , γ , and δ —represent the payoff the row player will receive when it plays strategy A or B with its partner, the column player, who can also play either strategy A or B .

Payoff Matrix 1. A two-player, symmetric game is as follows:

$$\begin{array}{cc} & \begin{array}{c} A & B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \begin{array}{cc} \alpha & \beta \\ \gamma & \delta \end{array} \end{array} \quad (1)$$

The Prisoner's Dilemma, in which $\gamma > \alpha > \delta > \beta$, is perhaps the most widely studied form of social interaction, having been used to model systems ranging from microorganisms (Conlin et al. 2014; [4], Frey 2010 [5]; West et al. 2006 [6]) to human societies (Axelrod and Hamilton [7]; Bowles and Gintis [8]). However, The Prisoner's Dilemma does not

accurately represent social dilemmas in which there is no strictly dominant strategy.

The Snowdrift Game (Sugden 1986 [9]), also known as the Hawk-Dove Game (Smith 1982 [2]), is a social dilemma in which an individual can participate (“cooperate”) or not (“defect”) in producing a public good. In this game, two individuals, each in her own car, want to get home, but their cars are blocked by a large pile of snow. For either to get home, at least one person must get out of her car to shovel the snow. But there is a cost to doing so, creating a strategic dilemma in which $\gamma > \alpha > \beta > \delta$. The game is relevant to a number of phenomena in biology, such as collective defense and resource extraction among microorganisms (Gore et al. 2009 [10]; Conlin et al. 2014 [4]), behavioral contests (Smith 1982 [2]), distributive justice in humans (Sugden 1986 [9]), behavioral diversification (Doebeli et al. 2004 [11]), and branching events (Wakano and Lehmann 2014 [12]).

In what we might think of as a “standard” model in evolutionary game theory, there is an infinitely large, well-mixed population within which individuals either cooperate or defect in one-time, pairwise games with each other. One generally assumes that individuals reproduce asexually and that the number of offspring each individual has is proportional to its fitness. This makes it easy to track how the frequencies of cooperators and defectors change in the population over time. The rate of change of a strategy is given by the replicator dynamics (Hofbauer and Sigmund 1998 [1]; Hofbauer et al. 1979 [13]; Taylor and Jonker 1978 [14]), an ordinary differential equation.

If a population of individuals is playing this game, then traditional evolutionary models, in which the dynamics are continuous and deterministic, predict a stable, interior equilibrium frequency of cooperators (Doebeli and Hauert 2005 [15]; Doebeli et al. 2004 [11]; Hauert and Doebeli 2004 [16]) (see (7)). Thus, the standard, deterministic model of the Snowdrift Game describes a scenario in which, even in the absence of facultative trait expression or heterozygote superiority, a stable polymorphism of behaviors can emerge.

However, all biological communities are finite, and many are small and organized into groups that together form a metapopulation (Gilpin 2012 [17]; Hanski 1999 [18]). This was certainly the case for ancestral hominins, which has plausibly influenced the evolution of human cooperation (Bowles and Gintis 2011 [8]), and it also appears to characterize many other taxa, including bacteria (Lieberman et al. 2016 [19]). It is therefore important to understand how a model with finite population size and metapopulation structure can change the predictions generated from models that assume a single population whose size tends to infinity.

Here we explore a discretization of the standard model of the Snowdrift Game. We consider a metapopulation composed of a finite set of discrete, nonintermixing groups, which are themselves composed of a finite set of discrete individuals who either cooperate or defect in the Snowdrift Game. The evolutionary dynamics both between and within groups are governed by a discrete time Moran process (Moran 1958 [20], 1962 [21]), a special case of a discrete time Markov chain.

Our analysis has two main results. First, we show that the combination of within-group stochasticity and group selection can promote the evolution of cooperation in the metapopulation and can even result in cooperation’s fixation.

This would be an impossible result were the evolutionary dynamics deterministic.

Second, we describe a phase transition for the fixation and extinction probabilities of cooperation in any finite group of a constant size whose members play the Snowdrift Game. Letting r stand for the cost-to-benefit ratio in a given Snowdrift Game, this threshold quantity, which we call r_* , is approximately equal to 0.796. If $r < r_*$, the probability cooperators will fix tends to 1 as group sizes go to infinity. If $r > r_*$, the probability cooperators will go extinct tends to 1 as group size goes to infinity. This is true so long as the starting frequency of cooperators is strictly between 0 and 1. (As we detail, more complicated dynamics occur when $r = r_*$.) The existence of this threshold quantity allows us to state a sufficient condition for cooperators to fix in a metapopulation. Moreover, while this threshold result comes about from taking the limit of group size, we in fact show that even when group size is fairly small (e.g., 100), the value of r effectively determines whether cooperators will fix or go extinct within a group. This threshold has no analogue in a deterministic model of the Snowdrift Game, and it makes the important differences between discrete and continuous models of evolution salient.

As we discuss, our results provide insight into the evolution of cooperation, particularly from a multilevel perspective (Luo 2014 [22]; Okasha 2006 [23]; Simon et al. 2013 [24]; Traulsen and Nowak 2006 [25]), evolving games (Akçay and Roughgarden 2011 [26]; Hashimoto and Kumagai 2003 [27]; Smead 2014 [28]), and the relationship between discrete and continuous models of evolutionary dynamics (Traulsen et al. 2005 [29]).

2. The Model

Here, we describe the within- and between-group Moran processes in our model.

Payoff Matrix 2 provides the payoff matrix for the two-player Snowdrift Game we will assume throughout this paper. Let b refer to the *benefit* of getting home and let c refer to the *cost* of shoveling snow, where $b > c$. Each driver can either get out of her car to shovel (C for “cooperate”) or stay in her car (D for “defect”). If both drivers cooperate, then each receives a net payoff of $b - c/2$, since both receive the benefit of going home, while the cost of shoveling is divided in half. If one driver cooperates and the other defects, then both get to go home, but the cooperator must pay the full cost of shoveling snow, receiving a net payoff of $b - c$, while the defector pays no cost, receiving a net payoff of b . If each driver stays in her car, neither pays the cost of shoveling, but neither gets to go home, so neither receives any reward. Following Zheng et al. (2007 [30]), we let r stand for the ratio of the cost of cooperating in the Snowdrift Game to the benefit of doing so ($r = c/b$). We can thereby speak of the “ r -value” of an instantiation of the game.

Payoff Matrix 2. The Snowdrift Game is as follows ($r = c/b$):

$$\begin{array}{cc}
 & \begin{matrix} C & D \end{matrix} \\
 \begin{matrix} C & \end{matrix} & \begin{matrix} b - \frac{c}{2} & b - c \end{matrix} \\
 \begin{matrix} D & \end{matrix} & \begin{matrix} b & 0 \end{matrix}
 \end{array} \tag{2}$$

Note that “cooperation” in this context refers to the shoveling snow behavior—that is, strategy C in Payoff Matrix 2. This strategy is not a form of altruism since cooperation can be in the interest of the actor, depending on the behavior of the other player. Strategy C in Payoff Matrix 2 coincides with a technical definition of cooperation (West et al. 2007 [31]): cooperative behaviors (i) increase the payoff to others and (ii) carry a benefit (or cost) to the actor contingent on the behavior of others.

2.1. Within-Group Moran Process. Suppose there is a metapopulation composed of a finite number n of discrete groups, indexed by j , so that $j \in \{1, 2, \dots, n\}$. The size of a given group j (N_j) is finite and constant, and each group in the metapopulation is of the same size. We assume the individuals in each group are hard-wired to either cooperate or defect in the Snowdrift Game. The evolutionary dynamics within each group are governed by a discrete time Moran process (Moran 1958 [20], 1962 [21]). At each time step, an individual in a group j is chosen to reproduce and have one offspring. The probability that a given individual is chosen to reproduce is proportional to its fitness relative to the average fitness of the individuals in its group. The behavior of an offspring is always identical to the behavior of its parent—that is, cooperators always beget cooperators, while defectors always beget defectors. At the moment it is born, an offspring replaces uniformly at random some individual in j , perhaps its parent, but not itself. The probability an individual will be replaced is unaffected by that individual’s phenotype or fitness and is always $1/N_j$. For our purposes here, we assume there is no migration between groups and no mutations during reproduction.

Since we are interested in modeling interactions that generate public goods that can be used by all, we will assume individuals play the Snowdrift Game with all of the members of their respective groups, including themselves, simultaneously. This is equivalent to using the expected fitness of random pairwise interactions among members of the group allowing for self-interaction.

Formally, we can represent the fitness of cooperators and defectors in a group j as follows. Let b_j and c_j index the parameters b and c for a given group j , let x_j stand for the frequency of cooperators in j , and let $1 - x_j$ stand for the frequency of defectors in j . The fitness of a cooperator and a defector in group j for a given value of x_j are given, respectively, by

$$w_{C,j} = \left(b_j - \frac{c_j}{2} \right) x_j + (b_j - c_j)(1 - x_j), \quad (3)$$

$$w_{D,j} = b_j x_j.$$

The average individual fitness of the members of a group j is given by

$$\bar{w}_j = x_j w_{C,j} + (1 - x_j) w_{D,j}. \quad (4)$$

The composition of a group can change in one of two ways: a cooperator can replace a defector, or a defector can replace a cooperator. Letting i_j stand for the number of

cooperators in a group j , we can represent the first transition as $(i_j \rightarrow i_j + 1)$ and the second as $(i_j \rightarrow i_j - 1)$. As described elsewhere (Fudenberg et al. 2006 [32]; Nowak 2006 [33]; Taylor et al. 2004 [34]), we can use the fitness given in (3)-(4) to calculate the probability, P , of each of these two transitions:

$$P(i_j \rightarrow i_j + 1) = \frac{w_{C,j}}{\bar{w}_j} \times \frac{i_j}{N_j} \times \frac{N_j - i_j}{N_j}, \quad (5)$$

$$P(i_j \rightarrow i_j - 1) = \frac{w_{D,j}}{\bar{w}_j} \times \frac{i_j}{N_j} \times \frac{N_j - i_j}{N_j}.$$

Since no other changes within a group are possible, the probability that the state of the group will not change is given by,

$$1 - P(i_j \rightarrow i_j + 1) - P(i_j \rightarrow i_j - 1). \quad (6)$$

Were each group well-mixed and infinitely large, each group j would converge to a stable, internal equilibrium frequency of cooperators (x_j^*), which is given by the following (Hauert and Doebeli 2004 [16]):

$$x_j^* = \frac{b_j - c_j}{b_j - c_j/2}. \quad (7)$$

However, because group size is finite in our model and there are no mutations, the only truly stable states of a group are the two absorbing states, in which cooperators fix ($x_j = 1$) or go extinct ($x_j = 0$). Nevertheless, the frequency of cooperators in a group in our model will often temporarily oscillate around what its internal equilibrium value would be were group size to be infinite; for a finite population, this is sometimes called its “quasi-equilibrium” (Shpak et al. 2013 [35]).

2.2. Between-Group Moran Process. Our model also involves a discrete time Moran process that occurs between groups. Within the metapopulation, a “parent” group is chosen to replicate, thereby producing a “daughter” group, which replaces uniformly at random some group in the metapopulation, perhaps its parent, but not itself. A daughter group of some group j that is created at time t has the same frequency of cooperators as j at t and has the same payoff matrix as j . (See Figure 1.)

The probability that a given group is chosen to reproduce is proportional to the average individual fitness of its members relative to the average individual fitness of the members of all the groups in the metapopulation. Taking the first and second derivatives of (4), we see that the average individual fitness of a group increases monotonically but nonlinearly as its frequency of cooperators increases:

$$\frac{d\bar{w}}{dx} = (c - 2b)x + (2b - c), \quad (8)$$

$$\frac{d^2\bar{w}}{dx^2} = c - 2b < 0.$$

In our model, a group’s fitness is equal to the probability it will give rise to a daughter group at the next time step.




FIGURE 1: *Two-Level Moran Process.* This is a condensed representation of the two-level, discrete time Moran process in our model. Each isolated cluster of red and blue ovals is a “group,” while the ovals themselves are “individuals,” which come in two varieties, cooperators (blue) and defectors (red). The set of six groups constitutes a metapopulation. Within a group, an individual is chosen to reproduce and replace some individual in the same group: the arrow represents an individual having an identical offspring, which replaces another individual in the group. (The replaced individual has a dashed border around it.) Across groups, a group is chosen to reproduce and replace some group in the metapopulation. The group in the bottom-left corner reproduces and replaces the group in the top-right corner. (The replaced group has a dashed border around it.)

Hence, a group’s fitness increases along with its frequency of cooperators. (One can of course consider other configurations of the relationship between individual fitness and group fitness—e.g., where a group’s fitness decreases as its frequency of cooperators increases—but we do not pursue such extensions here.)

In our model, we assume individual-level birth-death events occur more frequently than group-level reproduction-extinction events. Again, the abstract structure of the model itself does not require this. Presumably, the relative rates of the individual- and group-level events in a model should be governed by the target system one is modeling.

There is a straightforward way to calculate the probability that a given group j will replicate or die at the next time step. Letting $P_j(+1)$ and $P_j(-1)$ refer, respectively, to these two probabilities, we have

$$\begin{aligned} P_j(+1) &= \frac{\bar{w}_j}{\bar{w}_M} \times \frac{1}{n} \times \frac{n-1}{n}, \\ P_j(-1) &= \frac{\bar{w}'_j}{\bar{w}_M} \times \frac{1}{n} \times \frac{n-1}{n}, \end{aligned} \quad (9)$$

where \bar{w}_j is the average fitness of the individuals in group j , given by equation (4),

$$\bar{w}'_j = \frac{1}{n-1} \sum_{k:k \neq j} \bar{w}_k \quad (10)$$

refers to the average fitness of all groups other than group j , and \bar{w}_M is the average of the average individual fitness of each group in the metapopulation, which is given by

$$\bar{w}_M = \frac{1}{n} \sum_{j=1}^n \bar{w}_j. \quad (11)$$

The probability that there will be no change in the metapopulation is given by

$$1 - \sum_{j=1}^n [P_j(+1) + P_j(-1)]. \quad (12)$$

3. Methods

3.1. Simulations of Moran Process. We ran two classes of simulations. In the first class, we were concerned exclusively with the Moran process within a group. In the second class, we were concerned with our “complete” model, in which Moran processes occur both within and between groups. Both classes of simulations were implemented in R version 3.2.3 (R Core Team 2016 [36]). High Performance Computing (HPC) was necessary to perform some of the metapopulation simulations. This is because our code “counts” each individual in the metapopulation at each individual-level birth-death event and at each group-level replication-extinction event. This creates a computational bottleneck that can be nontrivial when group size is large ($N_j \geq 1000$). The source code is available by request.

3.1.1. Simulations of Within-Group Moran Process. In our first class of simulations, we sought to explore how the within-group evolutionary dynamics were affected by the size of a group and the group’s r -value. We therefore simulated the within-group evolutionary dynamics of nine groups, where each group was associated with a different r -value. The range of r -values that we simulated was $r \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$, where b was set to 10 and c ranged from 1 to 9 in increments of 1. We simulated a range of group sizes, $N_j \in \{10, 100, 1000, 10,000\}$, and we allowed each group to evolve for 100 generations, where 1 generation for a group of size N_j is equal to N birth-death events. The observed frequency of

cooperators within each group after a given number of birth-death events was plotted against the equilibrium frequency of cooperators predicted by the deterministic within-group model, as given by (7).

3.1.2. Simulations of Two-Level Moran Process. In our second class of simulations, we simulated the complete version of our model, in which there are discrete time Moran processes that occur both within groups and between groups.

In our first set of simulations within this class, we considered a scenario in which the groups in the metapopulation do not vary in their r -values. We seeded the metapopulation with 25 groups, where the frequency of cooperators in each was initially set to 0.10. We ran three versions of these simulations, which differed in the r -value we assigned to each group ($r = .25$; $r = .67$; $r = .90$). For each r -value, we ran simulations where $N_j \in \{10, 100, 1000\}$. We ran 500 independent simulations for each combination of r -value and group size. Within a given simulation, there were approximately 100 within-group generations and 20 between-group generations, where one generation for a metapopulation of size n is equal to n group replication events.

In our second set of simulations within this class, we considered a scenario in which the groups in the metapopulation initially varied in their r -values. To explore the sensitivity of the metapopulation dynamics to the r -values of each group, we ran two versions of this type of simulation. In the first version, we seeded the metapopulation with 25 groups and associated each with a unique r -value so that the average r -value was relatively low: the benefit b of cooperating was set to 10, while c ranged from 0.25 to 6.25, in increments of 0.25, resulting in an initial average r -value of 0.33. In the second version, we performed the same analysis, except that we associated each group with a unique r -value so that the average r -value was relatively high: b was set to 10, while c ranged from 7.79 to 9.94, in increments of approximately 0.09 on average, resulting in an initial average r -value in the metapopulation of 0.89.

These simulations were “truncated” in that they could stop before the metapopulation became monomorphic. We conducted truncated simulations because we were interested in documenting the evolutionary dynamics in the short and medium term, not only as one takes the limit of time. As our results show, the groups in the metapopulation, and the metapopulation itself, often spend a good deal of time at a polymorphic state. This is because the fixation times can be long, which underscores the need to not only consider the end state of the system. Eventually, this polymorphism will disappear, and so the simulations are an inaccurate representation of the dynamics of the system for an arbitrarily long sequence of events.

Comparison with Null Models. For each set of simulations within this class, we plotted the observed frequency of cooperators in the simulations against what the long-term, expected frequency of cooperators in the metapopulation ($\mathbb{E}[x_M]$) would be given two alternative, null models. In the first null model, there is a single, well-mixed population of individuals playing the Snowdrift Game, within which the

dynamics are deterministic. In the second, there is a metapopulation of groups whose members play the Snowdrift Game, where the within-group dynamics are deterministic, there is group reproduction, and there is no group selection (i.e., groups do not vary in their probability of replicating). Conveniently, the value for ($\mathbb{E}[x_M]$) for each of these null models is given by the same equation:

$$\mathbb{E}[x_M] = \frac{1}{n} \sum_{r \in R} n_r x_r^*, \quad (13)$$

where n_r represents the number of groups with a given r -value in the metapopulation, x_r^* stands for the equilibrium frequency of cooperators in a group of type- r , and the notation $r \in R$ indicates that we perform this operation for each r -value that is in the set of r -values that are represented in the metapopulation. In the case where there is a single, infinitely large population, we can think of this population as a “metapopulation” with a single, infinitely large group.

4. Results

4.1. Characterization of Within-Group Moran Process

4.1.1. Diffusion Approximation and Fixation Probability. Following Traulsen et al. (2005 [29]), we can obtain the diffusion approximation of the Moran process for a group whose members are playing a two-player game of the form given in Payoff Matrix 1. It is a diffusion process $X^{(N)} = (X_t^{(N)})_{t \geq 0}$ that solves the following SDE:

$$dX_t^{(N)} = \mu(X_t^{(N)}) dt + \frac{\sigma(X_t^{(N)})}{\sqrt{N}} dW_t, \quad (14)$$

where N is the size of the group (note that we have left out the j subscripts from N for ease of reading), W_t is the Wiener process (standard Brownian motion), and μ and σ are the “drift” and the “volatility” functions given, respectively, by

$$\begin{aligned} \mu(x) &= T^+(x) - T^-(x), \\ \sigma(x) &= \sqrt{T^+(x) + T^-(x)}, \end{aligned} \quad (15)$$

where

$$\begin{aligned} T^+(x) &= \frac{\alpha x + \beta(1-x)}{\bar{w}} x(1-x), \\ T^-(x) &= \frac{\gamma x + \delta(1-x)}{\bar{w}} x(1-x). \end{aligned} \quad (16)$$

Equation (14) describes how a change in the frequency of cooperators is dependent both on selection and on the stochasticity of the system that emerges because the population is finite. Observe that when the group’s size tends to infinity, the SDE reduces to an ODE:

$$\frac{dx(t)}{dt} = \mu(x(t)). \quad (17)$$

The precise connections between the Moran process and the corresponding SDE and the Moran process and the

corresponding ODE are given by the following lemma. A rigorous proof follows from standard stochastic analysis (Durrett 1996 [37]).

Lemma 1. *Let N be the total number of individuals and let $C_N(i)$ be the number of cooperators at the i th step in the Moran process. Suppose the initial fraction of cooperative players $C_N(0)/N$ tends to $x_0 \in [0, 1]$ as $N \rightarrow \infty$. Then, for all $\epsilon > 0$ and $T < \infty$, we have the following.*

(1) *Deterministic approximation: Moran process versus ODE:*

$$\lim_{N \rightarrow \infty} \mathbb{P} \left(\sup_{t \in [0, T]} \left| \frac{C_N([tN])}{N} - x(t) \right| < \epsilon \right) = 1, \quad (18)$$

where $x(t)$ is the solution to the ODE (17) starting with $x(0) = x_0$, “sup” is the supremum, and $[s]$ is the largest integer smaller than or equal to s .

(2) *Diffusion approximation: Moran process versus SDE: there exists a diffusion process $X_t^{(N)}$ defined on the same probability space as that of the Moran process, which starts at $X_0^{(N)} = x_0$, solves (14), and is such that*

$$\lim_{N \rightarrow \infty} \mathbb{P} \left(\sup_{t \in [0, T]} \left| \frac{C_N([tN])}{N} - X_t^{(N)} \right| < a_N \epsilon \right) = 1 \quad (19)$$

for any sequence a_N such that $\lim_{N \rightarrow \infty} (\log N / N a_N) = 0$.

Remark 2. Convergence (18) says that when the population size N is large, most sample paths of the process $C_N([tN])/N$ stay within distance ϵ from the solution $x(t)$ of the ODE for any time $t \in [0, T]$. That is, once we specify the error tolerance ϵ and the duration T , we can choose N large enough so that $|C_N([tN])/N - x(t)| < \epsilon$ uniformly in $t \in [0, T]$ with a probability close to one.

Remark 3. Convergence (19) says that $X_t^{(N)}$ is a better approximation than $x(t)$ and precisely quantifies the extent to which it is a better approximation. Specifically, (19) holds as long as $a_N \rightarrow 0$ slower than $\log N / N$. This roughly says that $\sup_{t \in [0, T]} |C_N([tN])/N - X_t^{(N)}| \rightarrow 0$ at a rate of order at least $O(\log N / N)$.

Remark 4. Lemma 1 remains true if we replace the discrete time Moran process by the continuous time version of the Moran process with rate N . This follows directly from Kurtz (1978 [38], theorems 2.2 and 3.3).

In evolutionary biology, one is often interested in the probabilities of fixation and extinction. Let $h(x)$ be the probability that the diffusion process $X_t^{(N)}$ hits 1 before 0 (i.e., cooperators fix), given that the process starts at $X_0^{(N)} = x$. In mathematical notation, we let $\tau := \inf\{t \geq 0 : X_t^{(N)} = 0 \text{ or } 1\}$ be the time it takes for the population to become monomorphic, and we let \mathbb{P}_x be the probability law of $X_t^{(N)}$ starting at $x \in [0, 1]$. Then,

$$h(x) := \mathbb{P}_x (X_\tau^{(N)} = 1). \quad (20)$$




FIGURE 2: The plot of $h(x)$. When $r = 0.3$, which is *below* r_* , the probability that cooperators will fix gets very large, regardless of the starting frequency of cooperators, as group size increases.

Note that $h(x)$ depends on N and the parameters in the game matrix.

This probability can be found in standard textbooks on the subject (e.g., Otto and Day 2007 [39], Chapter 15). We record the exact formula for our case in the following lemma.

Lemma 5. *Let $h(x)$ be the probability that cooperation fixes, defined in (20). Then,*

$$h(x) = \frac{\int_0^x e^{Nf(y)} dy}{\int_0^1 e^{Nf(y)} dy}, \quad (21)$$

where $f(y) = -2 \int_0^y (Cu + D)/(Au + B) du$ and

$$\begin{aligned} A &= \alpha + \gamma - \beta - \delta, \\ B &= \beta + \delta, \\ C &= \alpha - \beta - \gamma + \delta, \\ D &= \beta - \delta. \end{aligned} \quad (22)$$

4.1.2. Asymptotic Analysis Reveals a Threshold Cost-to-Benefit Ratio. Through an application of Laplace’s method to the explicit formula for the fixation probability obtained in Lemma 5, we discover a threshold quantity $r_* \approx 0.796$ such that when $r < r_*$ the probability cooperators will fix tends to 1 as group size goes to infinity (Figure 2). Likewise, when $r > r_*$, the probability cooperators will go extinct tends to 1 as group size goes to infinity (Figure 3). More complicated dynamics occur when r is roughly equal to r_* (Figure 4). Remarkably, this holds irrespective of the initial frequency of cooperators in the group. Moreover, the probability cooperation will fix or go extinct, given an arbitrary starting




FIGURE 3: The plot of $h(x)$. When $r = 0.9$, which is *above* r_* , the probability that cooperators will fix is negligible, regardless of the starting frequency of cooperators, as group size increases.




FIGURE 4: The plot of $h(x)$. When r is almost exactly equal to r_* , the probability that cooperators will fix is approximately $r_*/4 \approx 0.2$, regardless of the starting frequency of cooperators, as group size increases.

frequency of cooperators, can be extremely high even when group size is relatively small (e.g., $N_j = 100$) (Figures 2 and 3). To the best of our knowledge, this threshold quantity is unreported in the extant literature. We now describe its derivation in the “Snowdrift Game Threshold Theorem.” Note that this theorem applies only to the form of the Snowdrift Game presented in Payoff Matrix 2, not to the more general form of the game presented in Lemma 5.

Theorem 6 (snowdrift game threshold theorem). *There exists a threshold number $r_* \in (0, 1)$ such that, for all $x \in (0, 1)$, we have the limiting probability*

$$\lim_{N \rightarrow \infty} h(x) = \begin{cases} 0, & \text{if } r > r_*, \\ \frac{r_*}{4}, & \text{if } r = r_*, \\ 1, & \text{if } r < r_*, \end{cases} \quad (23)$$

where $h(x) := \mathbb{P}_x(X_\tau^{(N)} = 1)$ is the probability that cooperation goes to fixation, given that the starting point is $X_0^{(N)} = x \in [0, 1]$. Furthermore, $r_* \approx 0.79559250931$ is the unique solution of the following equation:

$$1 - \frac{r^2}{4} - 2(1-r) \ln\left(\frac{4-r}{2-2r}\right) = 0 \quad (24)$$

on the positive real axis.

Remark 7. The endpoints are given by $h(0) = 0$ and $h(1) = 1$ for all N .

Proof. Asymptotic analysis (see, e.g., Murray 1984 [40]) reveals that the large N behavior of the probability $h(x)$ is governed by the maximum of the function f . Recall the explicit formula and the notation in Lemma 5. For the Snowdrift Game, we have

$$\begin{aligned} f(y) &= -2 \int_0^y \frac{Cu + D}{Au + B} du \\ &= -2 \left(\frac{C}{A} y + \frac{AD - BC}{A^2} \ln\left(1 + \frac{A}{B} y\right) \right), \\ f'(y) &= -2 \frac{Cy + D}{Ay + B}, \\ f''(y) &= 2 \frac{AD - BC}{(Ay + B)^2}. \end{aligned} \quad (25)$$

Clearly, $f''(y) > 0$ and the maximum of f is attained at the endpoints $f(0)$ and $f(1)$. To derive (23), note that

$$\begin{aligned} f(1) - f(0) &= \frac{2b^2}{A^2} \theta(r) \\ \text{where } \theta(r) &= 1 - \frac{r^2}{4} - 2(1-r) \ln\left(\frac{4-r}{2-2r}\right). \end{aligned} \quad (26)$$

Elementary calculus shows that $f(0) < f(1)$ when $r \in [0, r_*]$, $f(0) = f(1)$ when $r = r_*$, and $f(0) > f(1)$ when $r \in (r_*, 1]$, where r_* is the solution of $\theta(r) = 0$ on the positive real axis, which was found to be $r_* \approx 0.79559250931$. When $f(0) = f(1)$ (i.e., $r = r_*$), we have, by the Laplace method (see, e.g., Murray 1984 [40]),

$$\begin{aligned} \lim_{N \rightarrow \infty} h(x) &= \left(\frac{1}{f'(0)} - \frac{1}{f'(1)} \right)^{-1} \frac{1}{f'(0)} = \frac{r_*}{4} \\ &\quad \forall x \in (0, 1). \end{aligned} \quad (27)$$

Result (23) then follows. \square

4.1.3. Simulations of Within-Group Moran Process. The results of our purely within-group simulations are presented in Figure 5. The jagged lines are the observed frequency of cooperators within a group with a particular r -value after a given number of birth-death events. The horizontal lines of the corresponding color are the equilibrium frequency of cooperators one would expect in that group were the group to be infinitely large, as given by (7).

In accord with analytic predictions, over a given number of generations, relatively smaller groups are more susceptible to larger fluctuations in group composition. Moreover, given a sufficient amount of time, and large enough groups, the within-group dynamics tend to follow a given pattern: in the short term, the frequency of cooperators in a group migrates toward the equilibrium frequency of cooperators given by (7), and in the medium term the frequency of cooperators in a group tends to oscillate around this equilibrium frequency, creating a quasi-equilibrium. This pattern becomes more evident as one moves from the simulations with the smallest group size (Figure 5(a)) to the simulations with the largest group size (Figure 5(d)).

4.2. Characterization of Metapopulation Dynamics. Theorem 6 allows us to state a sufficient condition for whether cooperators will fix (go extinct) in our two-level Moran model of the Snowdrift Game. As one takes the limit of group size, cooperation will fix (go extinct) in the metapopulation if each group in the metapopulation has an r -value below (above) r_* . Note that, in the limit of group size, the between-group selection process is irrelevant to the long-term dynamics of cooperation in the metapopulation. While this sufficient condition is a limit result, we stress that a group's size need not be very large for its r -value, relative to r_* , to substantially influence the probability cooperators will fix or go extinct.

When groups' r -values straddle r_* , when group sizes are small, or when one is concerned with evolutionary dynamics over shorter intervals of time, simulations are informative, owing to the challenge of formal analysis of these cases. Even though the two-level Moran process is a finite state Markov chain that counts the number of cooperators in each group, and so the probability of fixation can be computed explicitly, it is not obvious how to write down a closed-form expression in terms of (i) the number of groups, (ii) the number of individuals in each group, and (iii) the initial frequencies of cooperators in the groups when the number of groups and/or the size of the groups are as large as we are concerned with here.

The results of our metapopulation simulations are presented in Figures 6 and 7.

5. Discussion

5.1. Stochasticity and the Snowdrift Game. There is a good deal of work that explores how stochasticity affects the evolutionary dynamics of a population of individuals playing a game in general (Fudenberg et al. 2006 [32]; Nowak et al. 2004 [41]; Sandholm 2011 [42]; Traulsen and Hauert 2009 [43]). In our model, we explore a particular form of “demographic

stochasticity” (Rice 2008 [44]; Shaffer 1981 [45]; Shpak et al. [35]). Demographic stochasticity occurs when there is a difference between an individual's (or type's) expected fitness and its realized fitness. In our model, this occurs when the type (“cooperator” or “defector”) that is chosen to reproduce within a group is not the type with the higher fitness, and when the group that is chosen to replicate is not the group with the highest group fitness. Other forms of demographic stochasticity can emerge if an individual plays a random sample of the members of a finite population (Shpak et al. 2013 [35]), or if there are “shocks” to the payoffs of a game (Fudenberg and Harris 1992 [46]).

There is, in fact, a considerable amount of work that explores how stochasticity affects the evolution of cooperation in a single, unstructured population of individuals playing the Snowdrift Game (Ficici and Pollack 2000 [47]; Fogel et al. 1997 [48]; Fogel et al. 1998 [49]; Fudenberg et al. 2006 [32]; Shpak et al. 2013 [35]; Taylor et al. 2004 [34]; Voelkl 2010 [50]). While the results of these analyses are consistent with our mathematical and simulation results, to the best of our knowledge we are the first to report a threshold quantity for the fixation and extinction probabilities of a finite group of individuals playing the version of the Snowdrift Game presented in Payoff Matrix 2.

It is worth commenting on the general approach we have used to derive the value of this threshold quantity. In many traditional analyses in evolutionary game theory, one takes the limit of population size ($N \rightarrow \infty$) and then time ($t \rightarrow \infty$) to arrive at the long-term behavior of an evolutionary system. To derive the value of r_* , we have switched the order in which these limits are taken. We take ($t \rightarrow \infty$) to derive the formula for the fixation probability of cooperation. Then we take ($N \rightarrow \infty$) by applying Laplace's method (Murray 1984 [40]) to this formula. Determining the “correct” order in which limits are taken presumably depends on the nature of the system one has in mind. (For a discussion of this issue as it relates to mutations and population size, see Sandholm 2010 [51].)

Within the context of the Snowdrift Game, if the population is large and one is concerned with a short- or medium-term time horizon, then the evolutionary trajectory of the system may be adequately represented by a “deterministic approximation” of the microscopic dynamics (Sandholm 2011 [42]). Indeed, our simulation results show a marked correspondence to the deterministic dynamics after 100 generations when $N \geq 1000$ (Figures 5(c) and 5(d)). If, in contrast, one is concerned with smaller populations or the longer-term behavior of a larger population, the stochasticity of the system becomes increasingly important—see Figures 5(a) and 5(b) and Theorem 6, respectively.

5.2. Multilevel Selection and the Snowdrift Game. Our results provide insight into when within-group stochasticity and group selection do (and do not) result in evolutionary dynamics that are different from those in our two null models.

For instance, our results show that when the members of each group in the metapopulation play precisely the same parameterization of the Snowdrift Game, and when group




FIGURE 5: *Simulation results of within-group Moran process.* Each panel shows the evolution of cooperation in each of nine different groups over the course of 100 generations within each group. Within a panel, groups are all the same size but vary in their r -values. The range of r -values that were simulated was $r \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$, where b was set to 10 and c ranged from 1 to 9 in increments of 1. The jagged lines represent the frequency of cooperators in a group. As r -values get larger, the equilibrium frequency of cooperators gets smaller; thus, $r = 0.1$ is black, while $r = 0.9$ is teal. Each horizontal line of a corresponding color represents the equilibrium frequency of cooperators for that r -value given by (7).

size is small ($N = 10$), the combination of within-group stochasticity and group selection can promote the evolution of cooperation in the metapopulation above what its value would be given either of the two null models we consider (Figures 6(a)–6(c)).

Why does this occur? Within each group, the frequency of cooperators is oscillating. Because the probability that a group will replicate at the next group replication event is proportional to its frequency of cooperators at the time of the group replication event, those groups with a relatively higher frequency of cooperators during the group replication event are more likely to replicate. When this occurs, the frequency of cooperators in the metapopulation will “jump” upward, as in a jump diffusion process. Indeed, within-group stochasticity can result in cooperators fixing in a group, which

results in that group having the highest possible group fitness, making it possible that this group will then overtake the entire metapopulation. It is for this reason that we say within-group stochasticity in combination with group selection can promote the evolution of cooperation in the metapopulation, even its fixation, for a range of parameter values. (Note, as the number of groups in the metapopulation becomes increasingly large, the between-group selection dynamics become increasingly deterministic; thus, that there are a finite, rather small number of groups in our metapopulation simulations in fact dampens the efficacy of group selection; this modeling choice was intentional. We wanted to, if anything, load the die against the evolution of cooperation.)

However, our results also show that as the size of each group increases ($N \geq 100$), the within-group dynamics, over




FIGURE 6: *Two-level Moran model with no variation in r -values*. Each panel represents simulation results for our metapopulation model. Panels (a, d, g) show the results for $r = .25$; panels (b, e, h) show $r = .67$; panels (c, f, i) show $r = .9$. Each panel shows results for 100 independent simulations with populations of 10 (a, b, c), 100 (d, e, f), and 1000 (g, h, i). In each simulation, there were 25 groups with an initial cooperator frequency of .10, 500 group replication events, and 100 within-group generations. The red, black, and light blue (cyan) dots are the first, second, and third quartiles of the observed frequencies of cooperators in the metapopulation, while the thin black line is the mean frequency. The blue lines are the maximum and minimum observed frequency across 100 simulations. The dashed, horizontal line is the long-term, expected frequency of cooperation in the metapopulation assuming a null model given by (13). Note that the observed frequency of cooperators is substantially above its value predicted by the null models for smaller populations (a)–(c) and for one larger population with a high quasi-equilibrium value (d).

the time intervals we have simulated here, become increasingly deterministic. Over these time intervals, within-group stochasticity is less capable of resulting in jumps in the frequency of cooperators in the metapopulation after a group replication event, or resulting in cooperators fixing within a group (Figures 6(d)–6(h)). To be sure, because each metapopulation we simulated is composed of finite groups, the deterministic approximation of the dynamics will prove to be inaccurate in the long term. Our r_* results tell us that it is extremely likely cooperators would have overtaken the entire metapopulation in some of our simulations (Figures 6(a), 6(b), 6(d), 6(e), 6(g), and 6(h)) and gone extinct in others (Figures 6(c), 6(f), and 6(i)) had we allowed the simulations

to continue to evolve, owing to the distribution of r -values of the groups in those simulations.

When we allow the groups in the metapopulation to initially vary in their r -values, group selection can be quite efficacious, even for large group sizes (Figures 7(c)–7(f)). Regardless of group size, those groups with quasi-equilibria frequencies that are relatively closer to $x = 1$ will be more likely to replicate. When group sizes are small (e.g., $N = 10$), cooperators will fix in many of these groups (Figures 7(a) and 7(b)). But even as group size becomes larger ($N \geq 100$), and the within-group dynamics become increasingly deterministic, those groups with higher quasi-equilibria are more likely to replicate. For this reason, we again see the




FIGURE 7: *Two-level Moran model with variation in r -values*. Each panel represents simulation results for our model when groups in the metapopulation vary in their initial r -values. The panels (a, e, c) show the results for $r = .33$, while the panels (b, d, f) show $r = 0.89$. Each panel represents summary results for 100 independent simulations with populations of 10 (a, b), 100 (c, d), and 1000 (e, f) individuals. In each simulation, there were 25 groups with an initial cooperator frequency of .10, 500 group replication events, and 100 within-group generations. The red, black, and light blue (cyan) dots are the first, second, and third quartiles of the observed frequencies of cooperators in the metapopulation, respectively, while the thin black line is the mean frequency. The blue lines are the maximum and minimum observed frequency across 100 simulations. The dashed, horizontal line is the long-term, expected frequency of cooperation in the metapopulation assuming a null model given by (13). Note that the observed frequency of cooperators is greater than that predicted by the null models in all cases, sometimes substantially.

jump diffusion process of cooperators in the metapopulation that we saw in the version of our model in which each group initially had the same r -value. Within-group stochasticity and group selection are pushing the observed frequency of cooperators above what its expected frequency would be, given our null models.

Why have we considered a case in which we allow groups to vary in the parameter values of the Snowdrift Game their members play? First, in a biological context, if a metapopulation is composed of groups whose members are engaged in collective action that is *qualitatively* characterized by the Snowdrift Game, it would in fact be somewhat surprising if there were no *quantitative* differences between the payoff matrices associated with each group. (To be sure, there would also likely be variation in payoffs within groups, an interesting possibility we did not explore here.) For instance, these differences could emerge for reasons intrinsic to the group if, as an example, the members of some groups can more efficiently expend energy, so that the cost of participating in collective action is mitigated. Differences could also emerge because of ecological facts about the location of a group within the metapopulation, such as its proximity to energy sources or to external threats. To model the latter case, which we have not done, one would need to allow the payoff matrix to vary as a function of a group's location, which is an interesting avenue for future research.

Second, by allowing group payoff matrices to vary in our model, we learn something about the evolution of games themselves. Note that a central change brought about by selection and replacement of groups is a change to the payoff matrix of the underlying game. That is, as groups reproduce and go extinct, the payoffs of the game evolve. The possibility of an evolving game has been explored in other contexts, especially that of the Prisoners Dilemma, where it tends to promote cooperation in some contexts (Akçay and Roughgarden 2011 [26]; Worden and Levin 2007 [52]) and inhibit cooperation in others (Stewart and Plotkin 2014 [53]). Additionally, Hashimoto and Kumagai (2003 [27]) and Smead (2014 [28]) both present frameworks for the evolution of payoffs across a broad range of games. With respect to evolving games, our model is novel in two ways. First, we specifically target the Snowdrift Game and find that evolution of the payoffs promotes cooperation. Second, our model provides a mechanism for the evolution of the payoffs based on multi-level selection. This framework could be extended to a number of other games or settings for the purposes of modeling evolving payoffs.

Our model and results connect in illuminating ways to other work on multilevel selection in general, as well as to applications of multilevel selection to the Snowdrift Game in particular. Others have used a multilevel Moran model, or close variants, to explore the evolution of social traits (Hauert and Imhof 2012 [54]; Luo 2014 [22]; Simon et al. 2013 [24]; Traulsen and Nowak 2006 [25]). Notably, Traulsen and Nowak consider a model in which within-group dynamics are governed by the payoffs of playing a game, while in the metapopulation a group fissions (with some probability) when a group reaches a maximum size, replacing some other group in the metapopulation. Traulsen and Nowak provide

conditions for when cooperation will be “favored” in terms of a payoff matrix, the maximum size of a fissioning group, and the number of groups in the metapopulation. In their work, selection favors cooperation when a mutant cooperator has a higher probability of fixing than does a mutant defector. Traulsen and Nowak's results extend to cases in which the within-group dynamics are governed by a Snowdrift Game (SI Traulsen and Nowak 2006 [25]). (See Simon and Pilosov 2016 [55] for discussion of limitations of Traulsen and Nowak's model.)

Ours is a different model than Traulsen and Nowak's in that we consider a “strict” two-level Moran process, wherein both group size and metapopulation size are held constant. While two-level Moran models have been discussed before, and close variants of such a model have been described within the context of the Snowdrift Game (SI Traulsen and Nowak 2006 [25]), ours appears to be the most in-depth application of a two-level Moran model to the Snowdrift Game in the literature.

In fact, the dynamics of a strict two-level Moran model can differ in nontrivial ways from more general two-level Markov chain models. For instance, both Traulsen and Nowak (2006 [25]) and Simon et al. (2013 [24]) consider a model in which, when a group fissions, both the size and composition of the two resultant daughter groups are random variables. Because of this, a daughter group may end up with all cooperators (or none), even though its parent group was polymorphic. All else being equal, compared to a model, like ours, in which an entire group replicates, the stochasticity that emerges from the random seeding of daughter groups would appear to increase the probability that cooperation will fix in some groups and, hence, in the metapopulation as a whole. While such a model of group fissioning is biologically justifiable, in many cases more so than our own, our results show that one does not need the size or composition of daughter groups to be random variables in order for cooperation to evolve.

Finally, there has been other work that explores how group structure can affect the evolution of cooperation in the Snowdrift Game, where it has been shown to sometimes promote the evolution of cooperation, though not always (Hauert and Doebeli 2005 [15]; Killingback and Doebeli 1996 [56]). Hauert and Doebeli's work is particularly relevant to our own because they contrast the equilibrium of frequency of cooperators that would be found in a well-mixed population, given particular parameter values for a Snowdrift Game, with the frequency of cooperators observed in a population with individuals arranged on a lattice, given those same parameter values. The authors found that, for small x^* values, the equilibrium frequency of cooperators in a population with a lattice structure is larger than what is expected in a well-mixed population. However, for intermediate and large x^* values, the equilibrium frequency of cooperators found in the population is smaller than what is expected in a well-mixed population. They conclude, “unexpectedly, spatial structure reduces the proportion of cooperators for a wide range of parameters,” adding, “our results caution against the common belief that spatial structure is necessarily beneficial for cooperative behavior” (p. 643).

In contrast, our analysis shows that quite a different kind of population structure—namely, in which the groups are discrete units that are spatially isolated within a metapopulation and replicate as a function of their internal composition—can, in certain cases, facilitate the evolution of cooperation, rather than hinder it. Whether a lattice structure or the form of population structure we consider above is a better representation of a given biological system of course depends on the details of that system.

6. Conclusion

This article presented a model of discrete individuals packaged into discrete groups, where the individual- and group-level dynamics were stochastic. We showed that within-group stochasticity and group selection promote the evolution of cooperation in the metapopulation when group size is small, and also when group size is large and groups are allowed to vary in the payoff matrices their members play. We further showed that the long-term fate of cooperation as one takes the limit of group size is determined by the cost-to-benefit ratio of cooperating in the Snowdrift Game and that this ratio effectively determines the long-term fate of cooperation even when group size is fairly small (e.g., $N = 100$).

All biological populations are finite (though they can be quite large), and they are often nested within a larger metapopulation. When these factors are incorporated into a model that would otherwise assume deterministic dynamics and no group structure (or group structure but no group selection), the differences in the predicted levels of cooperation can be substantial.

Conflicts of Interest

The authors declare that they have no conflicts of interest regarding the publication of this paper.

Authors' Contributions

This study was led by Laurence Loewe, Brian McLoone, and Wai-Tong Louis Fan, all three of whom conceptualized the project and developed the model. The simulations were implemented and analyzed in *R* by Brian McLoone. Wai-Tong Louis Fan performed the main mathematical analyses, including the derivation of the value of r_* . Adam Pham and Rory Smead contributed to the mathematical development of the model and to the writing and organization of the manuscript. Rory Smead was responsible for the model's link to the literature on evolving games. The manuscript was written by Brian McLoone with input from all, particularly Rory Smead, Wai-Tong Louis Fan, and Laurence Loewe.

Acknowledgments

The authors would like to thank Caitlin Pepperell and Tracy Smith for discussion and for detailed comments on earlier versions of this paper and Neil Van Lysel and Christina Koch at the Center for High Throughput Computing at UW-Madison for help in conducting their simulations. Shishi Luo,

Burton Simon, and William Sandholm provided valuable comments on the manuscript. John Yin and David Schwartz raised important questions on the project in conversation. Brian McLoone acknowledges support from NSF Career Award 1149123 to Laurence Loewe and from an NHGRI training grant to the Genomic Sciences Training Program 5T32HG002760. Wai-Tong Louis Fan is supported by NSF Career Award DMS-1149312 to Sebastien Roch and by the Wisconsin Institute for Discovery at UW-Madison. Laurence Loewe is supported by NSF Career Award 1149123 and by the Wisconsin Institute for Discovery at UW-Madison.

References

- [1] J. Hofbauer and K. Sigmund, *Evolutionary Games and Population Dynamics*, Cambridge University Press, Cambridge, UK, 1998.
- [2] J. M. Smith, *Evolution and the Theory of Games*, Cambridge University Press, 1982.
- [3] J. M. Smith and G. R. Price, "The logic of animal conflict," *Nature*, vol. 246, no. 5427, pp. 15–18, 1973.
- [4] P. L. Conlin, J. R. Chandler, and B. Kerr, "Games of life and death: Antibiotic resistance and production through the lens of evolutionary game theory," *Current Opinion in Microbiology*, vol. 21, pp. 35–44, 2014.
- [5] E. Frey, "Evolutionary game theory: theoretical concepts and applications to microbial communities," *Physica A: Statistical Mechanics and its Applications*, vol. 389, no. 20, pp. 4265–4298, 2010.
- [6] S. A. West, A. S. Griffin, A. Gardner, and S. P. Duggle, "Social evolution theory for microorganisms," *Nature Reviews Microbiology*, vol. 4, no. 8, pp. 597–607, 2006.
- [7] R. Axelrod and W. D. Hamilton, "The evolution of cooperation," *American Association for the Advancement of Science: Science*, vol. 211, no. 4489, pp. 1390–1396, 1981.
- [8] S. Bowles and H. Gintis, *A Cooperative Species*, Princeton University Press, Princeton, 2011.
- [9] R. Sugden, "The economics of rights, cooperation, and welfare," *Blackwell*, 1986.
- [10] J. Gore, H. Youk, and A. Van Oudenaarden, "Snowdrift game dynamics and facultative cheating in yeast," *Nature*, vol. 459, no. 7244, pp. 253–256, 2009.
- [11] M. Doebeli, C. Hauert, and T. Killingback, "The evolutionary origin of cooperators and defectors," *Science*, vol. 306, no. 5697, pp. 859–862, 2004.
- [12] J. Y. Wakano and L. Lehmann, "Evolutionary branching in deme-structured populations," *Journal of Theoretical Biology*, vol. 351, pp. 83–95, 2014.
- [13] J. Hofbauer, P. Schuster, and K. Sigmund, "A note on evolutionary stable strategies and game dynamics," *Journal of Theoretical Biology*, vol. 81, no. 3, pp. 609–612, 1979.
- [14] P. D. Taylor and L. B. Jonker, "Evolutionarily stable strategies and game dynamics," *Mathematical Biosciences*, vol. 40, no. 1-2, pp. 145–156, 1978.
- [15] M. Doebeli and C. Hauert, "Models of cooperation based on the Prisoner's Dilemma and the Snowdrift game," *Ecology Letters*, vol. 8, no. 7, pp. 748–766, 2005.
- [16] C. Hauert and M. Doebeli, "Spatial structure often inhibits the evolution of cooperation in the snowdrift game," *Nature*, vol. 428, no. 6983, pp. 643–646, 2004.
- [17] M. Gilpin, *Metapopulation Dynamics: Empirical and Theoretical Investigations*, Academic Press, 2012.

- [18] I. Hanski, *Metapopulation Ecology*, Oxford University Press, 1999.
- [19] T. D. Lieberman, D. Wilson, R. Misra et al., "Genomic diversity in autopsy samples reveals within-host dissemination of HIV-associated *Mycobacterium tuberculosis*," *Nature Medicine*, vol. 22, no. 12, pp. 1470–1474, 2016.
- [20] P. A. P. Moran, "Random processes in genetics," *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 54, no. 1, pp. 60–71, 1958.
- [21] P. A. Moran, "The statistical processes in genetics," 1962.
- [22] S. Luo, "A unifying framework reveals key properties of multi-level selection," *Journal of Theoretical Biology*, vol. 341, pp. 41–52, 2014.
- [23] S. Okasha, *Evolution and the Levels of Selection*, Oxford University Press, 2006.
- [24] B. Simon, J. A. Fletcher, and M. Doebeli, "Towards a general theory of group selection," *Evolution*, vol. 67, no. 6, pp. 1561–1572, 2013.
- [25] A. Traulsen and M. A. Nowak, "Evolution of cooperation by multilevel selection," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 103, no. 29, pp. 10952–10955, 2006.
- [26] E. Akçay and J. Roughgarden, "The evolution of payoff matrices: Providing incentives to cooperate," *Proceedings of the Royal Society B Biological Science*, vol. 278, no. 1715, pp. 2198–2206, 2011.
- [27] T. Hashimoto and Y. Kumagai, "Meta-evolutionary Game Dynamics for Mathematical Modelling of Rules Dynamics," in *Advances in Artificial Life*, vol. 2801 of *Lecture Notes in Computer Science*, pp. 107–117, Springer Berlin Heidelberg, Berlin, Heidelberg, 2003.
- [28] R. Smead, "Evolving games and the social contract," *Complexity and the Human Experience*, Pan Stanford, Singapore, 61–80, 2014.
- [29] A. Traulsen, J. C. Claussen, and C. Hauert, "Coevolutionary dynamics: From finite to infinite populations," *Physical Review Letters*, vol. 95, no. 23, Article ID 238701, 2005.
- [30] D. F. Zheng, H. P. Yin, C. H. Chan, and P. M. Hui, "Cooperative behavior in a model of evolutionary snowdrift games with N-person interactions," *EPL (Europhysics Letters)*, vol. 80, no. 1, Article ID 18002, 2007.
- [31] S. A. West, A. S. Griffin, and A. Gardner, "Social semantics: Altruism, cooperation, mutualism, strong reciprocity and group selection," *Journal of Evolutionary Biology*, vol. 20, no. 2, pp. 415–432, 2007.
- [32] D. Fudenberg, M. A. Nowak, C. Taylor, and L. A. Imhof, "Evolutionary game dynamics in finite populations with strong selection and weak mutation," *Theoretical Population Biology*, vol. 70, no. 3, pp. 352–363, 2006.
- [33] M. A. Nowak, *Evolutionary Dynamics*, Harvard University Press, 2006.
- [34] C. Taylor, D. Fudenberg, A. Sasaki, and M. A. Nowak, "Evolutionary game dynamics in finite populations," *Bulletin of Mathematical Biology*, vol. 66, no. 6, pp. 1621–1644, 2004.
- [35] M. Shpak, S. H. Orzack, and E. Barany, "The influence of demographic stochasticity on evolutionary dynamics and stability," *Theoretical Population Biology*, vol. 88, pp. 47–56, 2013.
- [36] R Core Team, *R: A Language and Environment for Statistical Computing Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2016.
- [37] R. Durrett, *Stochastic calculus*, Probability and Stochastics Series, CRC Press, Boca Raton, FL, 1996.
- [38] T. G. Kurtz, "Strong approximation theorems for density dependent Markov chains," *Stochastic Processes and Their Applications*, vol. 6, no. 3, pp. 223–240, 1977/78.
- [39] S. P. Otto and T. Day, *A Biologist's Guide to Mathematical Modeling in Ecology and Evolution*, vol. 13, Princeton University Press, Princeton, NJ, USA, 2007.
- [40] J. D. Murray, *Asymptotic Analysis*, vol. 48 of *Applied Mathematical Sciences*, Springer-Verlag, New York, Second edition, 1984.
- [41] M. A. Nowak, A. Sasaki, C. Taylor, and D. Fudenberg, "Emergence of cooperation and evolutionary stability in finite populations," *Nature*, vol. 428, no. 6983, pp. 646–650, 2004.
- [42] W. H. Sandholm, "Stochastic evolutionary game dynamics: foundations, deterministic approximation, and equilibrium selection," in *Evolutionary game dynamics*, vol. 69 of *Proc. Sympos. Appl. Math.*, pp. 111–141, Amer. Math. Soc., Providence, RI, 2011.
- [43] A. Traulsen and C. Hauert, "Stochastic evolutionary game dynamics," in *Reviews of nonlinear dynamics and complexity*. Vol. 2, pp. 25–61, Wiley-VCH Verlag, Weinheim, 2009.
- [44] S. H. Rice, "A stochastic version of the Price equation reveals the interplay of deterministic and stochastic processes in evolution," *BMC Evolutionary Biology*, vol. 8, no. 1, article no. 262, 2008.
- [45] M. L. Shaffer, "Minimum Population Sizes for Species Conservation," *Bioscience*, vol. 31, no. 2, pp. 131–134, 1981.
- [46] D. Fudenberg and C. Harris, "Evolutionary dynamics with aggregate shocks," *Journal of Economic Theory*, vol. 57, no. 2, pp. 420–441, 1992.
- [47] S. Ficici and J. Pollack, "Effects of finite populations on evolutionary stable strategies," in *Proceedings of the 2000 Genetic and Evolutionary Computation Conference*, L. Darrell Whitley, Ed., pp. 927–934, Morgan-Kaufmann Publishers, San Francisco, CA, USA, 2000.
- [48] D. B. Fogel, G. B. Fogel, and P. C. Andrews, "On the instability of evolutionary stable strategies," *BioSystems*, vol. 44, no. 2, pp. 135–152, 1997.
- [49] G. B. Fogel, P. C. Andrews, and D. B. Fogel, "On the instability of evolutionary stable strategies in small populations," *Ecological Modelling*, vol. 109, no. 3, pp. 283–294, 1998.
- [50] B. Voelkl, "The 'hawk-dove' game and the speed of the evolutionary process in small heterogeneous populations," *Games*, vol. 1, no. 2, pp. 103–116, 2010.
- [51] W. H. Sandholm, "Orders of limits for stationary distributions, stochastic dominance, and stochastic stability," *Theoretical Economics*, vol. 5, no. 1, pp. 1–26, 2010.
- [52] L. Worden and S. A. Levin, "Evolutionary escape from the prisoner's dilemma," *Journal of Theoretical Biology*, vol. 245, no. 3, pp. 411–422, 2007.
- [53] A. J. Stewart and J. B. Plotkin, "Collapse of cooperation in evolving games," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 111, no. 49, pp. 17558–17563, 2014.
- [54] C. Hauert and L. A. Imhof, "Evolutionary games in deme structured, finite populations," *Journal of Theoretical Biology*, vol. 299, pp. 106–112, 2012.
- [55] B. Simon and M. Pilosov, "Group-level events are catalysts in the evolution of cooperation," *Journal of Theoretical Biology*, vol. 410, pp. 125–136, 2016.
- [56] T. Killingback and M. Doebeli, "Spatial evolutionary game theory: Hawks and Doves revisited," *Proceedings of the Royal Society B Biological Science*, vol. 263, no. 1374, pp. 1135–1144, 1996.

