Intriguing effects of selection intensity on the evolution of prosocial behaviors

Alex McAvoy[1,2] Andrew Rao[3] Christoph Hauert[4]

[1]Department of Mathematics, University of Pennsylvania, Philadelphia, PA 19104 USA
[2]Center for Mathematical Biology, University of Pennsylvania, Philadelphia, PA 19104 USA
[3]Harvard College, Cambridge, MA 02138 USA
[4]Department of Mathematics, University of British Columbia, Vancouver, BC V6T 1Z2 Canada

Abstract

In many models of evolving populations, genetic drift has an outsized role relative to natural selection, or vice versa. While there are many scenarios in which one of these two assumptions is reasonable, intermediate balances between these forces are also biologically relevant. In this study, we consider some natural axioms for modeling intermediate selection intensities, and we explore how to quantify the long-term evolutionary dynamics of such a process. To illustrate the sensitivity of evolutionary dynamics to drift and selection, we show that there can be a "sweet spot" for the balance of these two forces, with sufficient noise for rare mutants to become established and sufficient selection to subsequently spread. In particular, this balance allows prosocial traits to evolve in evolutionary models that were previously thought to be unconducive to the emergence and spread of altruistic behaviors. Furthermore, the effects of selection intensity on long-run evolutionary outcomes in these settings, such as when there is global competition for reproduction, can be highly non-monotonic. Although intermediate selection intensities (neither weak nor strong) are notoriously difficult to study analytically, they are often biologically relevant; and the results we report here suggest that they can elicit novel and rich dynamics in the evolution of prosocial behaviors.

1 Introduction

The study of prosocial behaviors, which entail personal costs in order to benefit others, has driven the development of many of the models and techniques of modern evolutionary game theory. In its simplest form, evolutionary game theory is a framework for modeling the population dynamics of two or more traits with frequency-dependent reproductive success. The localized interactions between competing types, which influence an individual's fecundity or survival, are called "games."

In the donation game [1], for example, each individual has one of two types: C ("cooperator/producer") or D ("defector/non-producer"). A producer pays a cost, c, to provide the opponent with a benefit, b. Non-producers do nothing. Two producers facing one another each receive $b - c$, whereas two non-producers each get nothing. But when paired with a producer, an individual is better off not making a donation, instead taking b while the opponent gets $-c$. Although the resulting conflict of interest is quite simple, prisoner's dilemma interactions like this game have been extremely influential in explaining the evolutionary origins of altruistic behaviors [2–4].

In order to understand their evolutionary dynamics, these kinds of prosocial traits are studied in conjunction with a rule for updating the population based on reproduction and death. Much of the theoretical work on the evolution of cooperation assumes weak selection, in which payoffs have small effects on reproductive rates. Weak selection is frequently a natural assumption for biological evolution because, in general, an organism with a particular mutation has myriad other traits, and so this one difference between those with wild-type phenotype and the organism with mutated phenotype is not likely to have large effects on overall reproduction rates [5, 6]. Similar remarks can be made about many cultural contexts. In a generic community of people, a single novel behavior might not be significantly noted or adopted, since there are many behaviors exhibited by each person, which may or may not be already shared with others, contributing to the social fabric.

In addition to being a biologically reasonable assumption in many situations, weak selection usually also facilitates the mathematical analysis of evolutionary models. While population structure, which constrains interactions and offspring dispersal, has long been understood (qualitatively) to influence the spread of prosocial behaviors like cooperation [7, 8], more recent studies have sought to uncover quantitative conditions for the spread of cooperation in populations with fine-grained spatial structure. Even in cases for which these conditions ultimately contain terms that can be interpreted as relatedness, the conditions themselves cannot readily be derived just from the intuition that population viscosity increases relatedness between interactants. As such, additional work is required, and weak selection is often a crucial component of such analyses. A well-known study of weak selection is that of death-Birth (*dB*) updating on large regular graphs, which resulted in the widely quoted "$b/c > k$" rule [9] that says cooperation evolves if the ratio of the benefit provided by a cooperator, b, to the cost of cooperation, c, exceeds the number of neighbors, k. (The capitalization in "death-Birth" refers to the step in which selection acts.) This work has been subsequently generalized in a number of ways, for instance to include finite population sizes [10, 11], distinct interaction and dispersal topologies [12], and arbitrary heterogeneous graphs [13].

On the other hand, there are other update rules that seem to never support cooperation, regardless of the population structure. The archetype of such processes is Birth-death (*Bd*) updating [14], which is inspired by the Moran process [15] and models global competition for reproduction. In well-mixed populations, where the literature historically begins, cooperation in prisoner's dilemma games is never favored [16]. For *Bd* updating, cooperation is also never favored on regular graphs like the cycle [17], in which every node has the same number of neighbors [9, 18–21].

The general question of whether prosocial behaviors like cooperation can be favored on any topology, under *Bd* updating and weak selection, is unresolved, but it is generally believed that this is not possible [9, 21, 22]. The intuition behind this is that in *Bd* updating, adjacent individuals, who are playing the prisoner's dilemma against each other, also compete to reproduce, since *Bd*

updating first selects a node to reproduce proportional to fecundity. Then, cooperators are directly helping adjacent and competing defectors to be selected to reproduce. Contrast this to *dB* updating, in which competition to be selected only occurs among the adjacent nodes to the one randomly chosen to die. Two such individuals might themselves not be neighbors, which lessens the chance that an altruist helps a competitor. Therefore, we would expect cooperation to be favored under *dB* updating but not under *Bd* updating [22]. This intuition is captured more broadly and succinctly by Débarre et al. [23], who conclude that "the evolution of social behaviour is determined by the scale at which social interactions affect competition."

However, weak selection is far from the only relevant paradigm in modeling evolutionary dynamics. There are also compelling examples of strong selection in biology. A classic case study is that of the peppered moth (*Biston betularia*) in England. Prior to the Industrial Revolution, the lighter-colored moths could effectively hide from predators camouflaged amongst lichen and tree bark. Once atmospheric pollution covered the environment in soot, darker peppered moths proliferated (and, in subsequent years, moths began to revert back to their original color) [24–28]. Changing environmental conditions amplified the influence of one particular trait, moth color, above other potential mutations. Rapid temperature changes, e.g. due to climate change, have also been observed to result in increased selection pressures, related to the ability to bear heat, in populations of trees (*Fagus sylvatica*) [29] and wild birds (*Parus major*) [30]. Thus, even for populations in which weak selection might be a viable assumption under normal circumstances, environmental changes can (unsurprisingly) result in non-weak selection for some traits.

Compared with the regime of weak selection, less theoretical work has been done to understand social evolution under intermediate and strong selection, but these areas are not entirely unexplored. "Strong" selection frequently means that the selection intensity parameter is highest, essentially the polar opposite of weak selection. For prisoner's dilemma games on cycles, under *Bd* updating and strong selection, it is possible that a community of cooperators can resist invasion and fixation of a mutant defector [31]. More generally, intermediate and strong selection have been studied for multi-player games on the cycle [32], where it is shown that both the multi-player stag hunt game [33] and the snowdrift game [34] can favor cooperation, regardless of the intensity of selection.

In this study, we first discuss several issues that become increasingly important when modeling stronger selection in evolutionary games. These include how to convert payoff to fecundity, how mutants arise within the population, and how to quantify the effects of selection on evolutionary dynamics. We then highlight two important points using examples in simple structured populations. The first demonstrates that models with stochastic payoffs can exhibit completely different dynamics than their deterministic (mean-field) analogues. The second finding is that selection, of the proper intensity, can favor the evolution of prosocial behaviors under *Bd* updating, which challenges the consensus that *Bd* updating is necessarily a hostile setting for altruism. In particular, prosocial behaviors can evolve even when there is global competition for reproduction.

2 Evolutionary games and the intensity of selection

When selection is neither weak nor strong, three concerns immediately arise regarding how to model and evaluate evolutionary dynamics. The first issue relates to how payoffs accrue in games are converted to fecundity. A number of functional forms have been proposed in the literature, but all of them are effectively equivalent under weak selection, so there is no need to distinguish between different payoff-to-fecundity maps in that case. When selection is not weak, we show that there is a unique payoff-to-fecundity map, provided some natural qualitative assumptions hold.

The second concern is with how mutants arise in the population. Here, we focus on processes in which there is either no mutation or for which mutation is sufficiently rare. As a result, once a mutant does appear, the evolutionary process will drive the mutant type either to extinction or fixation prior to the arrival of another mutant. While different kinds of mutant appearance distributions are often considered in the literature on weak selection, here the issue becomes a bit more nuanced.

Finally, for processes with sufficiently rare mutation, we arrive at the question of how to quantify the evolutionary success of a trait. Under weak selection, this question is typically approached in one of two ways, either absolutely ("does selection increase the fixation probability of cooperators?") or relatively ("is the fixation probability of cooperators greater than that of defectors?"). These two conditions are, of course, still valid and reasonable for non-weak selection. However, outside of weak selection, there are two additional, analogous conditions that use a stronger selection intensity as a reference point. Under weak selection, these additional conditions collapse to the previous two conditions.

2.1 Payoff accrual

Here, we consider payoffs arising from simple prosocial behaviors. In particular, matrix games among two players do not specify how individuals accrue payoffs through interactions in a population. Naturally, this represents a prerequisite for understanding and analyzing the resulting evolutionary dynamics.

In fact, there is no canonical way to define such a quantity, particularly when the population structure is heterogeneous. Two of the most popular ways of defining an individual's overall payoff are through accumulation and averaging [35]. In the former, an individual simply adds together the payoffs from all pairwise encounters, while in the latter these payoffs are averaged. Both methods involve deterministic payoffs, in which each individual receives an amount that is uniquely specified by the population's structure and configuration of producers and non-producers. However, both kinds of payoff aggregation in a two-player game are just the beginning of the story of prosocial traits in structured populations. More specifically, the donation game is additive [36–38], which means that an individual's payoff in an interaction can be decomposed into the payoff from his or her action plus the payoff from the action of the opponent. Therefore, rather than thinking of payoffs as the result of pairwise *interactions*, in the donation game it is also natural to think of payoffs arising from a sequence of individual *actions*, one for each individual in the population. Incidentally, this action-based interpretation is why we use the terms "producer" and "non-producer" in place of the more traditional labels "cooperator" and "defector."

This interpretation naturally gives rise to other kinds of payoff accounting, such as concentrated or diffuse benefits [39]. When benefits are concentrated, every producer chooses a random neighbor and pays a cost, c, to confer a benefit, b, on this neighbor only ("*cf*-goods," where "*c*" indicates that the benefit is concentrated and "*f*" indicates that the cost is fixed). No action is taken to benefit other neighbors. When benefits are diffuse, each producer pays the same cost, c, but divides the benefit equally among all neighbors so that each neighbor gets at least a small benefit ("*ff*-goods" since both the total benefit and the total cost are fixed, irrespective of the number of neighbors a producer has). The latter may be viewed as a mean-field version of the former since, under concentrated benefits, the expected payoff to any given neighbor is exactly what this neighbor would receive under diffuse benefits. The classical setting is that of "*pp*-goods," wherein the total benefit and the total cost are proportional to the number of neighbors a producer has. Depictions of these traits, together with the classical payoff scheme, are shown in **Fig. 1**. The structures considered here, the cycle (**Fig. 1***(a)*) and the star (**Fig. 1***(b)*), represent two extremes with regards to payoff accrual. In the former, payoffs are distributed in the same way throughout the population; in the latter, there is a single individual whose connections vastly outnumber those of others.

2.2 Converting payoff to fecundity

In evolutionary games, the payoff, u, to an individual is converted to relative fecundity through a payoff-to-fecundity map, $F(u)$. We impose the following natural and intuitive restrictions on this map. *(i)* In a mechanistic approach to evolutionary dynamics, birth and death should be treated separately rather than amalgamating them into one quantity known as "fitness" [40]. Each of these separate components make sense only when non-negative; and as such, the propensity to reproduce, i.e. fecundity, should also be non-negative. *(ii)* Increasing payoffs are associated to greater abilities in competition, and hence F should be non-decreasing. *(iii)* Continuity of F is natural in the sense that it implies sufficiently small changes in payoffs do not result in large jumps in fecundity. *(iv)* Selection intensity is interpreted as an amplification of the payoffs prior to the conversion to fecundity. *(v)* Finally, competition should be independent of any baseline payoff or wealth shared by everyone. In particular, this last assumption allows one to focus on the effects of a particular game on competition to reproduce. Thus, in slightly more detail, we have the following assumptions:

(i) $F(u) \geqslant 0$ for every $u \in \mathbb{R}$;

(ii) F is non-decreasing;

(iii) F is continuous;

(iv) selection intensity scales payoffs, i.e. the fecundity of payoff u at intensity $\delta \geqslant 0$ is $F(\delta u)$;

(v) for every competition within a group $I \subseteq \{1, \ldots, N\}$, the probability that $i \in I$ is chosen for reproduction is invariant under adding a constant, K, to the payoffs of all competing individuals in I. That is, if u_i and F_i are the payoff and fecundity of individual i, then

$$\frac{F_i(u_i)}{\sum_{j \in I} F_j(u_j)} = \frac{F_i(u_i + K)}{\sum_{j \in I} F_j(u_j + K)}. \tag{1}$$

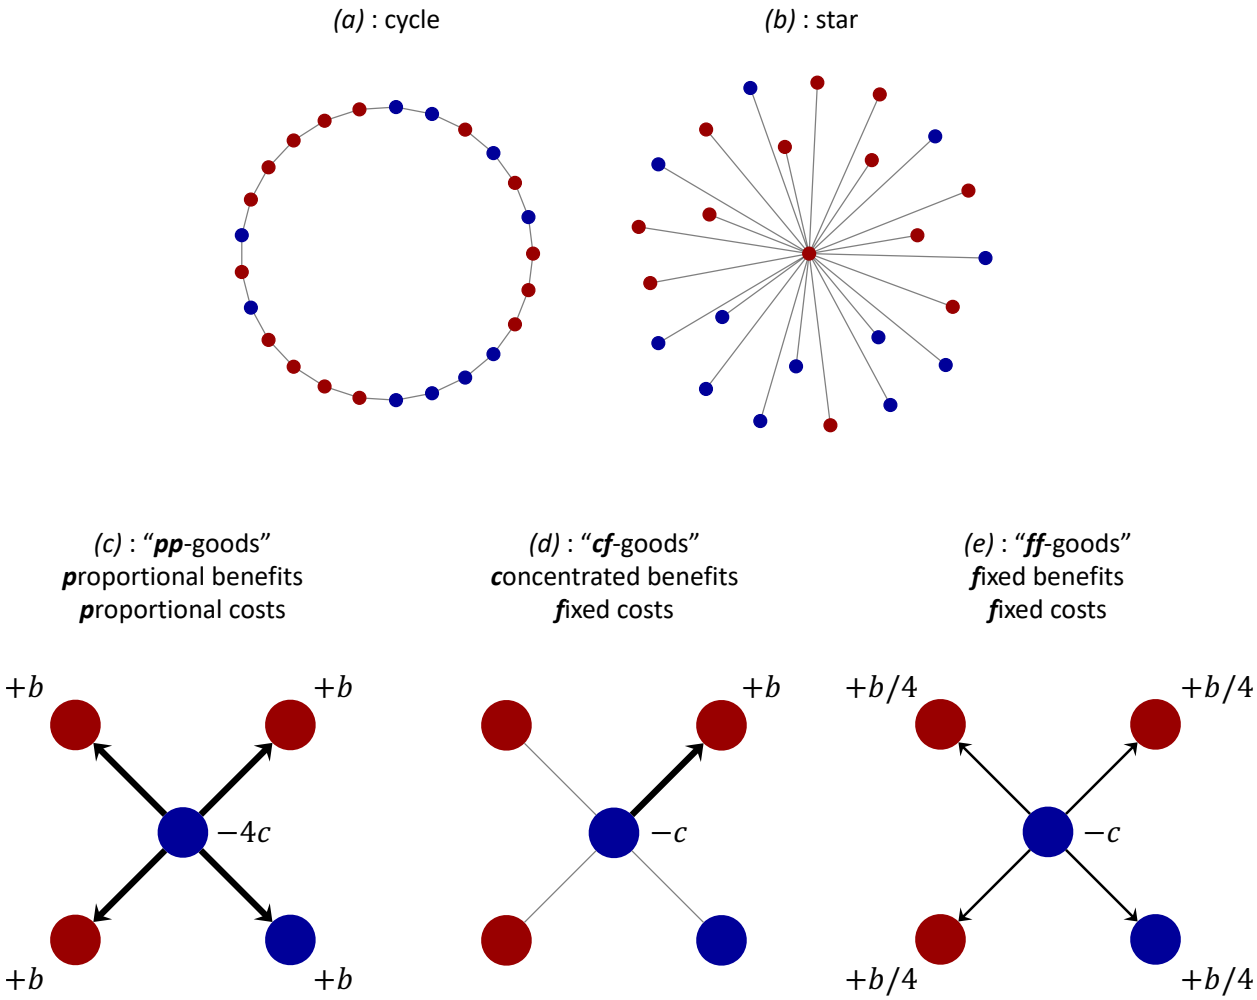

Figure 1: *(a,b)* Two examples of spatial structure in a population of size $N = 25$. Each node indicates an individual, who is either a producer (blue) or non-producer (red). Links between nodes represent possible interaction partners and dispersal patterns. The graph depicted in *(a)* is a cycle, which is a periodic structure in which every individual has exactly two neighbors. The graph in *(b)* is a star, a heterogeneous structure consisting of one central "hub" and $N - 1$ "leaf" nodes at the periphery. *(c,d,e)* Prosocial behaviors derived from the production of social goods. *(c)* represents the typical model, wherein each producer donates b to each neighbor, at a per-neighbor cost of c. In *(d)*, every producer confers the entire benefit produced, b, on a single, randomly-chosen neighbor. In *(e)*, this benefit is divided up evenly among all k neighbors. In both *(d)* and *(e)*, every producer in the population generates a total benefit of b and pays a total cost of c; the difference between the two is the way in which this benefit is distributed. Payoffs are a stochastic function of state in *(d)* and a deterministic function of state in *(e)*. Whereas *(c)* is *interaction*-based, the models of concentrated and diffuse benefits in *(d)* and *(e)* are more naturally seen as *action*-based payoff specifications.

In the Appendix (Section A.1), we show that, up to a rescaling of the selection intensity, these assumptions lead to a unique payoff-to-fecundity map, $F(u) = e^{\delta u}$ for selection intensity δ. Exponential payoff-to-fecundity maps like this one are already commonly used in the literature [32, 41–45], although often with little discussion for the choice of this particular functional form. One argument is that it satisfies assumption *(i)* for all possible payoffs, u, precluding complications arising from functional forms like $F(u) = 1 + \delta u$ [21, 46] or $F(u) = 1 - \delta + \delta u$ [16], which can both be negative when u is sufficiently small. Under weak selection ($\delta \ll 1$), these functional forms are all equivalent, so one can simply make an arbitrary choice [35]. Another argument is that this kind of exponential mapping matches imitation processes inspired by statistical mechanics, such as a standard model of pairwise comparison dynamics [47, 48]. Here, we see that this mapping arises as a consequence of several natural assumptions on the effects of payoffs on competition.

If selection acts on survival rather than on reproduction, we obtain a similar result. However, assumptions *(ii)* and *(v)* need to be modified accordingly. We discuss selection acting on survival further in the Appendix (Section A.2).

2.3 Fixation probabilities and mutant initialization

A standard method for determining the effects of selection on a trait is through fixation probability [16, 49, 50]. Let $\rho_{C,i}$ (resp. $\rho_{D,i}$) denote the probability that a single producer (resp. non-producer) at location i generates a lineage that takes over the population. Suppose that $\mu_{C,i}$ (resp. $\mu_{D,i}$) is the probability that a mutant C (resp. D) appears at location i in the monomorphic all-nonproducer (resp. all-producer) state. The mean fixation probabilities of the two types are then

$$\rho_C^{\mu} := \sum_{i=1}^{N} \mu_{C,i}\rho_{C,i}; \tag{2a}$$

$$\rho_D^{\mu} := \sum_{i=1}^{N} \mu_{D,i}\rho_{D,i}. \tag{2b}$$

Under uniform initialization, the probability that an invading mutant (of either type) appears at location i is $1/N$. Uniform mutant appearance is relevant, for example, when mutations arise in the adult generation of a population (e.g. adults who are equally likely to explore new behaviors or ideas) [51], or when death rates are roughly the same throughout the population. External mutagens, such as cosmic rays [52–54], can also introduce mutations more uniformly within a population. We denote the mutant-appearance distribution, μ, under uniform initialization by μ^{unif}, which satisfies $\mu_{C,i}^{\text{unif}} = \mu_{D,i}^{\text{unif}} = 1/N$. In this case, μ^{unif} is independent of the selection intensity, δ. We denote by ρ_C^{unif} and ρ_D^{unif} the corresponding mean fixation probabilities.

Another natural distribution is based on mutations that arise during reproduction, and it depends on selection intensity since these mutations are influenced by the rates at which nodes are replaced. In the monomorphic all-producer (resp. all-non-producer) state, let $d_{C,i}$ (resp. $d_{D,i}$) be the probability that the individual at location i is replaced by the offspring of another individual, in one step of the process. Since mutations arise during reproduction, higher replacement rates

naturally correspond to greater probabilities of mutant appearance. In the limit of rare mutation, we obtain what is known as "temperature" initialization [14, 55], in which the probability that a mutant arises at location i is $\mu_{C,i}^{\text{temp}} = d_{D,i} / \sum_{j=1}^{N} d_{D,j}$ (resp. $\mu_{D,i}^{\text{temp}} = d_{C,i} / \sum_{j=1}^{N} d_{C,j}$) following the all-non-producer (resp. all-producer) state. We denote by ρ_C^{temp} and ρ_D^{temp} the mean fixation probabilities under temperature initialization.

Notably, uniform and temperature initialization need not coincide. As an example, we consider Birth-death updating, in which an individual is first selected to reproduce with probability proportional to fecundity, and then the offspring replaces a random neighbor. Let $\left(w_{ij}\right)_{i,j=1}^{N}$ be the adjacency matrix for the population structure, and let $p_{ij} := w_{ij} / \sum_{k=1}^{N} w_{ik}$ be the probability of moving from i to j in one step of a random walk on the graph. If F_i denotes the fecundity of individual i at a given point in time, then the probability that node i is replaced is

$$d_i := \sum_{j=1}^{N} \frac{F_j}{\sum_{k=1}^{N} F_k} p_{ji}. \tag{3}$$

Since $\sum_{i=1}^{N} d_{C,i} = \sum_{i=1}^{N} d_{D,i} = 1$, we have $\mu_{C,i}^{\text{temp}} = d_{D,i}$ and $\mu_{D,i}^{\text{temp}} = d_{C,i}$.

Consider the star graph (**Fig. 1***(b)*). In the monomorphic non-producer state, everyone has the same payoff, regardless of selection intensity. As a result, the probability that the hub reproduces is $1/N$, and the probability that a leaf node reproduces is $1 - 1/N$. This means that a mutant arises at a leaf with probability $1/N$ and at the hub with probability $1 - 1/N$. In the monomorphic producer state, reproduction rates depend on the selection intensity. When $\delta \ll 1$, a mutant arises at a leaf with probability approximately $1/N$ and at the hub with probability approximately $1 - 1/N$. But as δ grows, the former probability approaches 1, while the latter approaches 0. **Fig. 2** illustrates temperature initialization on the star in these two states for producers of *ff*-goods (**Fig. 1***(e)*) under Birth-death updating.

2.4 Quantifying the effects of selection

There are four natural ways to quantify selection favoring C:

$\left(A^0\right)$ $\rho_C(\delta) > \rho_C(0)$ (absolute; in comparison to neutral drift);

$\left(A'\right)$ $\rho_C'(\delta) > 0$ (absolute; in comparison to selection intensity δ);

$\left(R^0\right)$ $\rho_C(\delta) > \rho_D(\delta)$ (relative to D; in comparison to neutral drift);

$\left(R'\right)$ $\rho_C'(\delta)\,\rho_D(\delta) > \rho_C(\delta)\,\rho_D'(\delta)$ (relative to D; in comparison to selection intensity δ).

Condition $\left(A^0\right)$ disregards the fixation probability of non-producers and requires selection to increase the fixation probability of producers over its value at neutral drift. The metric is "absolute" in the sense that it doesn't take into account the fixation probability of D. Outside of weak selection, this condition involves the comparison of two processes that are quite different.

Condition $\left(A'\right)$ states that, at a selection intensity of δ, increasing the selection intensity tends to improve the chances that a rare producer can invade and take over a population of non-producers.

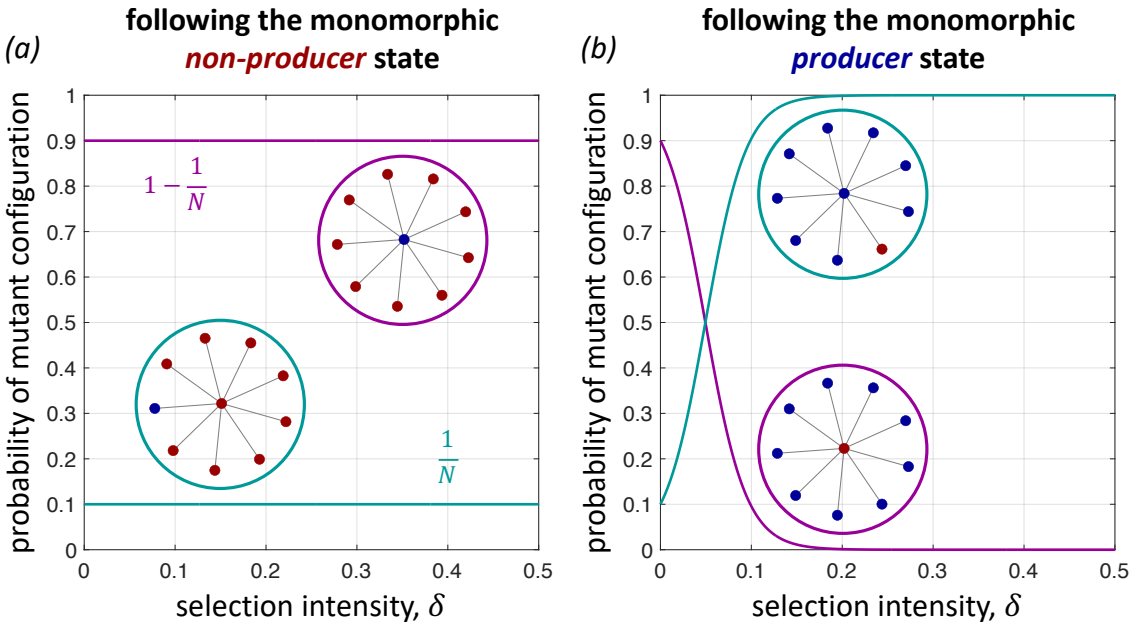

Figure 2: Temperature-based mutant initialization for *ff*-goods under Birth-death updating. *(a)* Following the monomorphic non-producer state, a mutant producer appears either in the hub (purple) or a leaf (turquoise). Since everyone has the same payoff (zero) in the all-non-producer state, the death rates are the same as those of neutral drift, for all $\delta \geqslant 0$. With probability $1/N$, the hub gives birth and the mutant offspring is propagated to a leaf node. Otherwise, with probability $1 - 1/N$, a leaf node gives birth and the mutant arises in the hub. *(b)* In contrast, death rates in the monomorphic producer state depend on selection intensity. Outside of weak selection, the hub is overwhelmingly more likely to reproduce than a leaf node, meaning a mutant appears in a leaf (turquoise). Although a non-producer at the hub, surrounded by producers, would have a high payoff and thus a high probability of spreading its behavior throughout the population, this mutant configuration appears with probability rapidly approaching zero as δ grows (purple). Parameters: $N = 10$, $b = 5$, and $c = 1$.

This condition is "absolute" for the same reason $\left(A^0\right)$ is, but the point of comparison is the fixation probability of C at δ, and δ is not necessarily small.

Condition $\left(R^0\right)$ says that C is favored relative to D if the probability that a rare producer takes over a population of non-producers exceeds the probability that a rare non-producer takes over a population of producers. Equivalently, this condition says that

$$\frac{\rho_C(\delta)}{\rho_C(\delta)+\rho_D(\delta)} > \frac{\rho_C(0)}{\rho_C(0)+\rho_D(0)} = \frac{1}{2}. \tag{4}$$

In a process with sufficiently rare mutation, the quotient $\rho_C/\left(\rho_C+\rho_D\right)$ represents the amount of time spent in the all-producer state; the remaining time is spent in the all-non-producer state [56, 57]. Thus, condition $\left(R^0\right)$ is a measure of C relative to D, and its point of comparison is neutral drift. If this condition holds, then at a selection intensity of δ, the population spends a larger fraction of its time in the all-producer state than it does under neutral drift. This kind of measure, based on invasion and counter-invasion, is a natural way to quantify evolutionary stability in finite populations [58, 59].

Finally, condition $\left(R'\right)$ is equivalent to $\left(\rho_C(\delta)/\left(\rho_C(\delta)+\rho_D(\delta)\right)\right)' > 0$. It says that, at a selection intensity of δ, increasing the selection intensity tends to increase the fraction of time the population spends in the all-producer state. Thus, it provides the same kind of relative measure as $\left(R^0\right)$, but the point of comparison is δ.

In the limit of weak selection, $\delta \ll 1$, conditions $\left(A^0\right)$ and $\left(A'\right)$ are equivalent, as are conditions $\left(R^0\right)$ and $\left(R'\right)$. Conditions $\left(A^0\right)$ and $\left(A'\right)$ can be equivalent to $\left(R^0\right)$ and $\left(R'\right)$ under weak selection, for example in games on graphs with the "equal gains from switching" property such as the donation game [9, 16, 36], but they need not always be [see 60]. For stronger selection, all four conditions are generally distinct. In the subsequent examples, we consider these conditions as a way to classify and understand the effects of selection on the evolution of producers of social goods.

3 Examples

We begin with an example of death-Birth (*dB*) updating, an evolutionary update rule known to support the evolution of prosocial behaviors in structured populations [9–11, 13, 19]. Under *dB* updating, an individual is first selected for death uniformly-at-random from the population. The neighbors of this individual then compete to fill the vacancy with probability proportional to relative fecundity. A neighbor's offspring inherits the parental type (either C or D), replaces the deceased individual, and the process repeats. *dB* updating is homogeneous in the sense that the individual chosen for replacement is always chosen uniformly, regardless of the population's structure or selection intensity. In particular, uniform and temperature initialization coincide, so there is no need to consider the two mutant-appearance distributions separately.

Fig. 3 depicts the evolutionary dynamics of three kinds of producers of social goods (Section 2.1) on the cycle (**Fig. 1***(a)*). Due to the homogeneity of the cycle, all initially-rare mutant configurations are equivalent. Since every individual on the cycle has exactly two neighbors, the

difference between *pp*-goods (**Fig. 1***(c)*), those of benefit and cost proportional to the number of neighbors, and *ff*-goods (**Fig. 1***(e)*), those of fixed benefit and cost, amounts to a rescaling of the selection intensity. Hence, **Fig. 3***(c)* appears as simply a horizontally-stretched version of **Fig. 3***(a)*.

What is more notable is the fact that *cf*-goods (**Fig. 3***(b)*), those of concentrated benefit and fixed cost, and its mean-field analogue, *ff*-goods (**Fig. 3***(c)*), have completely different behavior apart from when the selection parameter is small. Suppose that a non-producer arises in a population of producers of *ff*-goods. In order for non-producers to fix, the population must transition through a state containing a cluster of exactly two non-producers. In order to acquire additional non-producers from this state, a producer on the boundary must be selected for death. The non-producer competing to fill the vacancy has a payoff of $b/2$, while the producer on other side of the deceased has a payoff of $b-c$. As a result, the probability of acquiring an additional non-producer in this state is $(2/N)\left(e^{\delta b/2}/\left(e^{\delta b/2}+e^{\delta(b-c)}\right)\right)$, where the factor of 2 is due to there being two producer/non-producer boundaries. Provided $b>2c$, this probability approaches zero as δ grows, which effectively precludes non-producers from replacing producers. But consider what happens in this state for producers of cf-goods. If the competing non-producer gets b, then the competing producer necessarily has at most $b-c$ because the shared producer, who died, can donate to exactly one of these two individuals. Thus, the probability of acquiring an additional non-producer does not vanish in this state as δ grows. Similar reasoning in the other transient states explains the dynamics of non-producers depicted in **Fig. 3***(b)*.

The similarities between all three panels in **Fig. 3** when δ is sufficiently small is no coincidence. In the limit of weak selection, it is known that one can replace stochastic payoffs with deterministic payoffs (i.e. expected payoffs) without a loss of generality [39]. In particular, *cf*- and *ff*-goods exhibit identical dynamics when selection is sufficiently weak, on any population structure. **Fig. 3** demonstrates that this finding does not extend to stronger selection intensities. If one were to model stochastic payoffs using an expected-payoff approach, the model would show that stronger selection favors producers relative to non-producers when, in actuality, that is not the case (**Fig. 3***(b)*). A description of how the exact numerical calculations in **Fig. 3** were carried out may be found in the Appendix (Section A.3).

While *dB* updating highlights important differences between the dynamics stochastic and deterministic payoffs, much more curious phenomena arise under Birth-death (*Bd*) updating. Under *Bd* updating, in each time step an individual is chosen for reproduction with probability proportional to relative fecundity. The clonal offspring then replaces a random neighbor. One major difference between *dB* and *Bd* updating is that fecundity acts locally in the former and globally in the latter. As a result, *Bd* updating is not homogeneous when selection acts on the traits; differential birth rates result in neighbors of highly successful individuals being replaced more than those with lower fecundity.

For both kinds of mutant initialization, uniform and temperature, **Fig. 4** shows the evolutionary dynamics of three kinds of social goods on the star (**Fig. 1***(b)*). For *pp*-goods, producers are disfavored, both absolutely and relative to non-producers, for all δ when mutant initialization is uniform. When mutant initialization is temperature-based, producers are still disfavored absolutely, but they are favored relative to non-producers provided δ is sufficiently large. The reason for this difference is precisely the behavior described in **Fig. 2***(b)*, wherein the only state with a rare non-

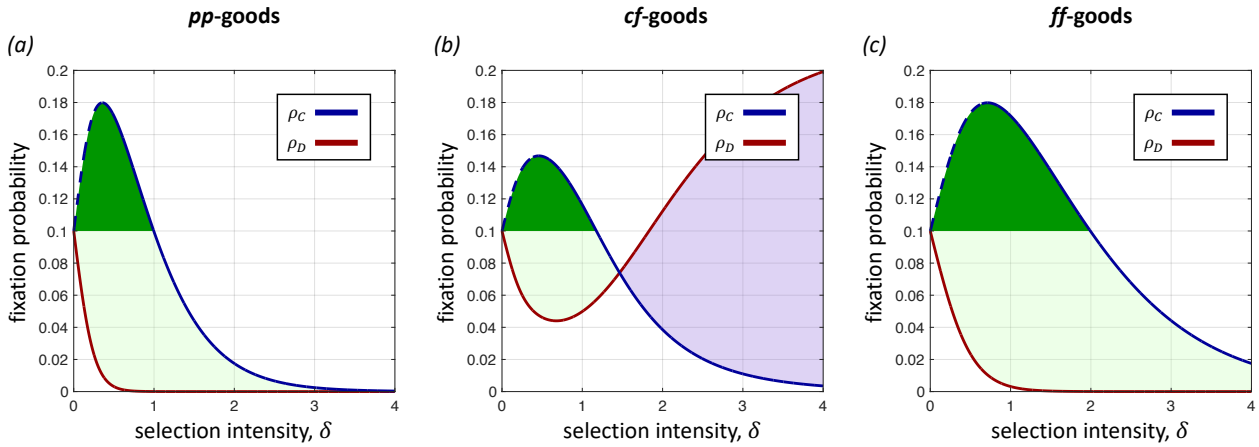

Figure 3: Fixation probabilities with uniform initialization (= temperature initialization) on the cycle of size $N = 10$ under *dB* updating (exact calculations). Since the cycle is a regular graph, the difference between *pp*-goods and *ff*-goods amounts to a scaling of the selection intensity. In both cases, *(a)* and *(c)*, C is always favored by selection relative to D (dark+light green), i.e. condition $\left(R^0\right)$. But in *(b)*, when benefits are concentrated, ρ_D eventually exceeds ρ_C; beyond this intensity of selection, C is disfavored relative to D (purple, condition $\left(R^0\right)$). Notably, non-weak selection highlights significant qualitative differences between the stochastic payoff scheme of *(b)* and its mean-field analogue, *(c)*. In all three cases, C is favored absolutely for small δ, both with respect to neutral drift (dark green, condition $\left(A^0\right)$) and with respect to δ (dashed line, condition (A')). For simplicity, we do not illustrate condition (R'). Parameters: $b = 4$ and $c = 1$.

producer becomes vanishingly unlikely to appear as an initial condition. We note that under weak selection, the behavior depicted in all panels of **Fig. 4** is essentially the same, with producers disfavored both absolutely and relative to D.

More fascinating is the fact that under both *cf*- and *ff*-goods, for both kinds of mutant initialization, producers can be favored both absolutely and relative to non-producers, provided selection is neither too weak nor too strong (**Fig. 4***(b,c,e,f)*). Weak selection initially favors non-producers and disfavors producers, consistent with traditional observations about *Bd* updating. However, past a certain selection intensity, fixation probabilities of producers start to increase, even going significantly above their respective values at neutrality. Those of non-producers, in contrast, decline and then dip below their respective values at neutrality. Once δ becomes sufficiently large, the fixation probabilities of producers then fall below their neutral values once again. In the case of uniform initialization, the fixation probability of non-producers eventually rises above that of producers (**Fig. 4***(b,c)*).

What is happening at these intermediate selection intensities? For one, drift is still strong enough that the population can transition into a favorable state for producers, such as the one with a producer at the hub and a small number of producers at the periphery. At this point, using *ff*-goods as an example, the producer at the hub has a greater payoff than all other individuals in the population due to the fact that they pay a single cost, c, but receive the whole benefit, b, from each of several neighbors. While drift is strong enough to visit such a state, so too is selection strong enough to drive the population to an all-producer state. The individual at the hub is chosen for reproduction disproportionately often, and the offspring must replace individuals at the periphery, eventually driving non-producers to extinction. When selection is too weak, these

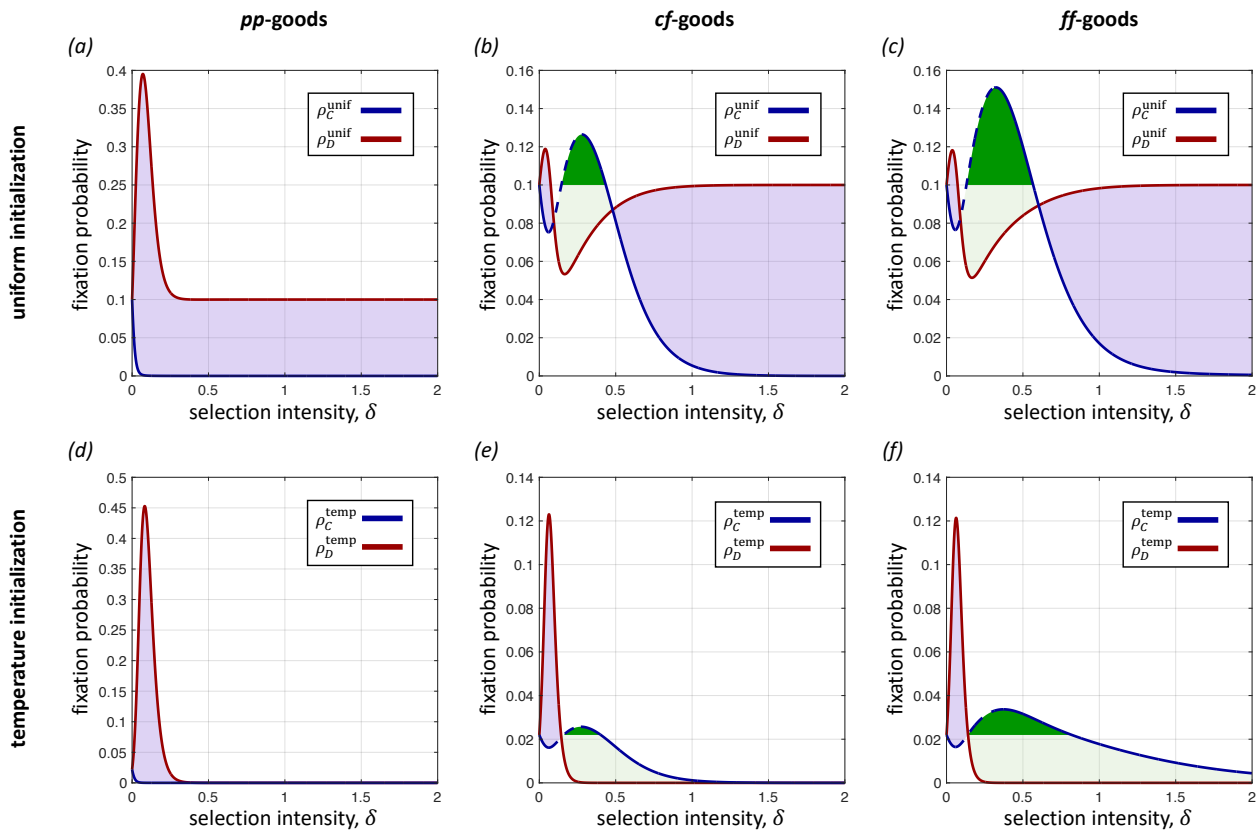

Figure 4: Fixation probabilities on the star of size $N = 10$ under *Bd* updating (exact calculations). *(a,b,c)* depict the dynamics of social goods under uniform initialization, whereas *(d,e,f)* show those of death-rate-based temperature initialization. In each panel, we see that weak selection disfavors producers, consistent with earlier studies. However, for both kinds of mutant initialization, we observe the remarkable fact that producers of *cf*- and *ff*-goods can be favored, both absolutely (dark green, condition (A^0); dashed line, condition (A')) and relative to non-producers (light+dark green, condition (R^0)), for a range of intermediate selection intensities. This range represents a "sweet spot" for the population, wherein drift is strong enough to allow for configurations favorable to producers to get established, and selection is strong enough to push those favored configurations toward fixation. The purple regions represent selection intensities for which C is disfavored relative to D (condition (R^0)). For simplicity, we do not show condition (R'). Parameters: $b = 5$ and $c = 1$.

favorable states are not favored enough to overcome the effects of drift. When selection is too strong, these favorable states are unlikely to be visited in the first place.

For *pp*-goods, we do not observe the same dynamics at intermediate selection intensities. Once a mutant producer occupies the hub, it is comparatively more difficult for producers to replace non-producers. In contrast to *ff*-goods, there are additional costs to a producer of *pp*-goods at the hub since it must pay a separate cost for each neighbor. Moreover, *pp*-goods offer additional benefits for those at the leaves since the central producer provides each neighbor with a full (rather than fractional) benefit. We note that, due to the rapid decay of mutant appearance at the hub in the all-producer state (**Fig. 2***(b)*) under temperature initialization, producers are eventually favored relative to non-producers, even though neither one is favored absolutely at strong selection intensities (**Fig. 4***(d)*).

Continuing with this example, we also consider significantly larger population sizes. **Fig. 5** depicts the fixation probabilities of producers of *ff*-goods under temperature initialization for population sizes ranging from $N = 10$ to $N = 2000$, extending the example of **Fig. 4***(f)*. At its peak, strikingly, the fixation probability of a single producer is nearly 140-fold its neutral value when $N = 2000$ (**Fig. 5***(a)*). Further nonlinearities also arise as the population size increases. For sufficiently large N, ρ_C^{temp} has two distinct local maxima as a function of δ, the first of which being attributed to the leaf nodes and the second being attributed to the hub (**Fig. 5***(b)*).

In **Fig. 5**, the peak for the fixation probability of mutants arising at a leaf node comes before that of mutants arising at the hub. For both kinds of initial mutant configurations, the population must transition to a state with a producer at the hub and a single producer at a leaf node on the way to fixation. If the mutant starts at a leaf node, then this producer has payoff $-c$ and competes to reproduce with $N-2$ individuals with payoff 0 and one individual with payoff b. If the mutant starts at the hub, then this mutant has payoff $-c$ and competes with $N-1$ individuals with payoff $b/\left(N-1\right)$. In both cases, transitioning to this intermediate two-mutant state is disfavored by selection, but this transition is less likely when the mutant arises in a leaf than when it arises in the hub. Therefore, selection must be weaker in the former case than in the latter for the two transitions to have equal probability. Finally, the presence of a "minor" peak in **Fig. 5***(b)* can be explained by the fact that, after transitioning from a producer at the hub to one at the hub and one at the leaf, there is a non-zero probability of losing a mutant again, which results in a single producer at a leaf.

4 Discussion

The four ways in which we have quantified selection favoring type C, namely $\left(A^0\right)$, $\left(A'\right)$, $\left(R^0\right)$, and $\left(R'\right)$ in Section 2.4, are all long-term measures [50, 61] of the evolutionary dynamics. As long as the selection intensity is finite, the population must eventually reach a monomorphic state, so all of the fixation probabilities considered here are well-defined. However, one might argue that the transient dynamics can be more important in models of non-weak selection than in those with weak selection. For example, when there is global competition to reproduce, stronger selection can essentially localize the dynamics for long periods of time. To see why, consider a highly-connected hub within a larger, heterogeneous graph with *Bd* updating. Locally, this hub resembles the star of **Fig. 1***(b)*. If the hub and all of its neighbors are producers, then the hub has an extremely large pay-

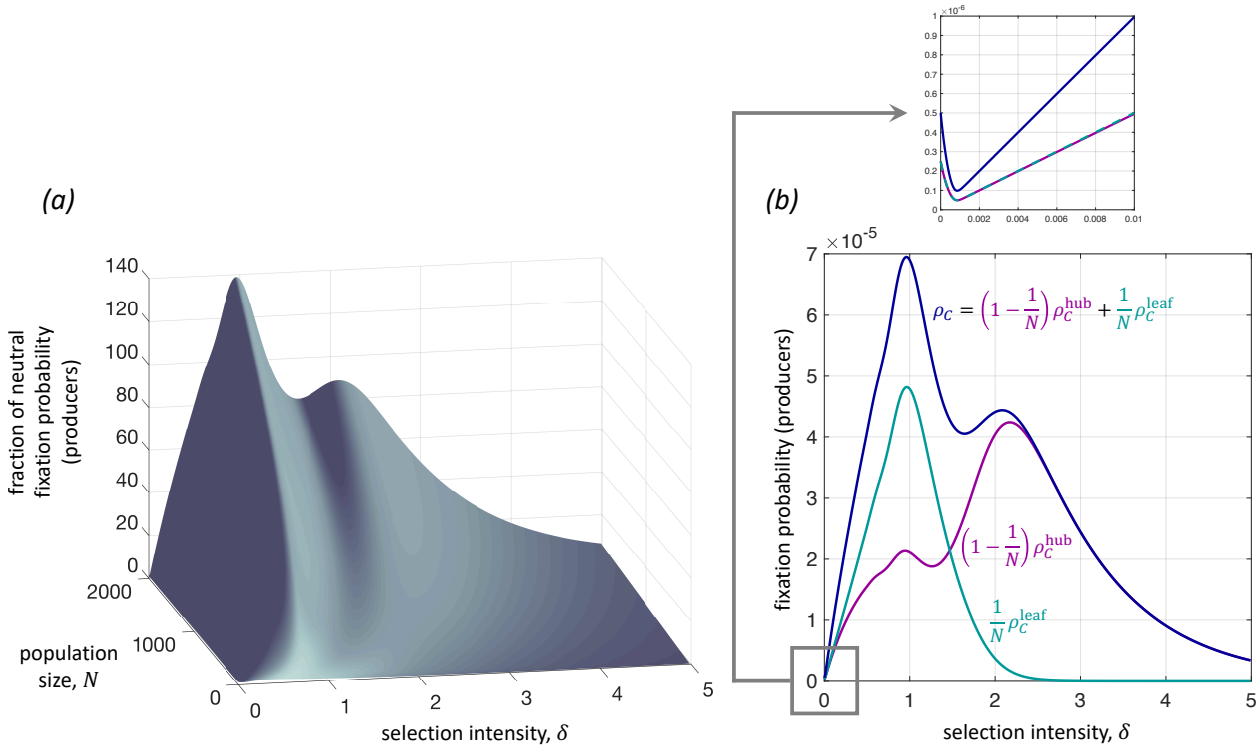

Figure 5: Fixation probability of producers of *ff*-goods on the star, as a function of size and selection intensity (exact calculations). These curves generalize the blue curve for the fixation probability of producers in **Fig. 4***(f)*. *(a)* Unlike in the case of smaller population sizes, ρ_C can have multiple local maxima when N is larger. *(b)* When $N = 2000$, the peak fixation probability of producers is magnified by a factor of nearly 140 over its neutral value. For the star, one peak can be attributed to mutant producers arising at the leaf nodes (ρ_C^{leaf}), while the other can be attributed to mutant producers arising at the hub (ρ_C^{hub}). For all population sizes depicted here, producers are initially disfavored by weak selection (sufficiently small δ; inset). Parameters: $b = 5$ and $c = 1$.

off relative to other individuals in the population. Since competition is global, this hub is involved in all tournaments for reproduction. If at least some other individuals in the population are non-producers, an individual other than the hub must be selected to reproduce for the state to change. But this might happen only rarely, especially when fecundity grows exponentially in payoff and selection intensity. These transient states, which resemble the metastable coexistence states observed in other games [62], are also relevant in addition to the long-term metrics considered here.

The conditions imposed on the payoff-to-fecundity map in Section 2.2 do not necessarily need to hold in every population. This is especially true of *(v)*, which says that shared background payoff/wealth does not affect competition to reproduce. It might be reasonable that fecundity is zero whenever u is less than some threshold payoff, in which case competition is not purely relative. Fecundity could also be a concave or linear function of payoff. All of these functional forms depend on the underlying nature of competition for reproduction and/or survival. The assumption that we ignore any background payoff shared by all members of a population is reasonable in light of the focus on just two strategies, C and D. After all, if D is the status quo, which includes all background traits irrelevant to the adoption of a new trait, C, and fecundity could be interpreted as being associated to the adoption of C only, rather than to the organism as a whole. What is more important in this model is that the exponential payoff-to-fecundity is not arbitrary. It arises as the unique map satisfying a number of simple and intuitive conditions based on the nature of competition. Ideally, other payoff-to-fecundity maps should be derived from similar assumptions and not chosen arbitrarily from some class of maps (e.g. from the class of concave maps).

To our knowledge, our examples provide the first demonstrations that selection can favor the evolution of prosocial traits under Birth-death updating. It is worth mentioning that *Bd* updating and weak selection can support cooperation in weaker social dilemmas. For example, in the snowdrift game, spatial structure can increase the equilibrium frequency of cooperation from that of an unstructured population, and indeed the right choice of payoffs can eliminate defectors altogether [18]. The sharp contrast between *Bd* and *dB* updating seen in the prisoner's dilemma does not hold for the snowdrift game, which may be attributed (in part) to the fact that defection is not a dominant action in this game [22]. In more stringent conflicts of interest like the production of social goods, where defection is dominant, non-weak selection can draw richer behavior out of social dilemmas, even resulting in the proliferation of altruistic traits when there is global competition to reproduce.

This proliferation of producers, which is a result of a "sweet spot" for the balance of drift and selection, is closely related to how heterogeneous populations generate wealth stratification [63]. While *pp*-goods can result in an unequal distribution of payoffs, this inequality is significantly amplified by *ff*-goods [39]. On the star, for precisely this reason, states that are favorable to the hub for *pp*-goods are even more favorable for *ff*-goods. As a result, the overlap between the selection intensities resulting in the appearance of favorable configurations and the selection intensities driving those configurations toward the all-producer state is larger for *ff*-goods than it is for *pp*-goods. In fact, there is no such overlap at all for *pp*-goods, as shown in **Fig. 4***(a,d)*.

Our use of exact calculations in small populations with simple spatial topologies is motivated by two factors. First, there are technical obstacles, computationally, to determining whether selection favors cooperation if the intensity of selection is not sufficiently small [64]. Second, one way of accounting for payoffs here (concentrated benefits) involves stochastic payoffs. These are

considerably more difficult to analyze directly, although they are perhaps even more relevant than deterministic payoffs in real populations. And, importantly, they cannot be modeled using expected (mean) payoffs unless one can prove that such a model results in the same dynamics. More than anything, both the update rules and the population structures should be understood as simple examples to illustrate what we believe is important behavior that can arise in the dynamics of models with non-weak selection. The patterns explored here through the relevant abstractions of the cycle (or regular graphs more generally) and the star provide important reference points for the dynamics in general networks. We anticipate that teasing out the effects of non-weak selection in other heterogeneous populations will be difficult, but that ultimately these kinds of models will exhibit much more fascinating behavior than those with extremely weak or strong selection.

Acknowledgments

We are grateful to Christian Hilbe, Yoichiro Mori, Joshua Plotkin, and Qi Su for helpful conversations. A.M. was supported by the Simons Foundation (Math+X Grant to the University of Pennsylvania). C.H. acknowledges funding by the National Science and Engineering Research Council Canada (NSERC), grant RGPIN-2015-05795.

Appendix

A.1 Uniqueness of the payoff-to-fecundity map

Here, we show that under the five assumptions, *(i)–(v)*, outlined in Section 2.2, it must be true that $F(u) = e^{\alpha u+\beta}$ for some $\alpha \geqslant 0$ and $\beta \in \mathbb{R}$. To see why, suppose first that the group in which competition for reproduction takes place is $I = \{1,2\}$. In this case, condition *(v)* is equivalent to the equation

$$\frac{F(u_1+K)}{F(u_2+K)} = \frac{F(u_1)}{F(u_2)} \tag{A.1}$$

holding for every $u_1, u_2, K \in \mathbb{R}$. We note that condition *(i)*, which says that $F(u) \geqslant 0$ for every $u \in \mathbb{R}$, implies that $F(u) > 0$ for every $u \in \mathbb{R}$ when combined with condition *(v)*. Otherwise, competition for reproduction, defined by condition *(v)*, would not be well-defined. Letting $u_2 = 0$, we then have $F(u_1+K) = F(u_1)F(K)/F(0)$ for every $u_1, K \in \mathbb{R}$. Since F is also continuous by *(iii)*, it is well-known that this functional equation implies that $F(u) = e^{\alpha u}/F(0)$ for some $\alpha \in \mathbb{R}$ [65]. F being a non-decreasing function of u, condition *(ii)*, is equivalent to $\alpha \geqslant 0$. Letting $\beta = -\log F(0)$, we then obtain $F(u) = e^{\alpha u+\beta}$. This function clearly also satisfies all of the desired properties for larger sets, I, which gives the claimed result.

Since selection intensity scales payoffs, *(iv)*, the fecundity attributed to payoff u at intensity δ is $e^{\alpha\delta u+\beta}$. But competition is relative, so the probability that i reproduces when competing with the individuals in I is $F_i/\sum_{j\in I} F_j = e^{\alpha\delta u_i}/\sum_{j\in I} e^{\alpha\delta u_j}$ (also called a "softmax function" [66] or a "Gibbs measure" [67]). Thus, without a loss of generality, we may assume that $\beta = 0$. As a result,

up to a rescaling of the selection intensity by α, there is a unique payoff-to-fecundity function satisfying the five assumptions imposed in the text.

A.2 Fecundity versus survival

If selection acts on survival rather than on reproduction, then the assumptions in Section 2.2 can be modified in a straightforward manner. Specifically, if F quantifies how likely an individual is to be chosen for death, then the corresponding assumptions are

(i) $F(u) \geqslant 0$ for every $u \in \mathbb{R}$;

(ii) F is non-increasing;

(iii) F is continuous;

(iv) selection intensity scales payoffs, i.e. the survival quantifier associated to payoff u at intensity $\delta \geqslant 0$ is $F(\delta u)$;

(v) for every competition within a group $I \subseteq \{1, \ldots, N\}$, the probability that $i \in I$ is chosen for death is invariant under adding a constant, K, to the payoffs of all competing individuals in I. That is, if u_i and F_i are the payoff and survival quantifier of individual i, then

$$\frac{F_i(u_i)}{\sum_{j \in I} F_j(u_j)} = \frac{F_i(u_i + K)}{\sum_{j \in I} F_j(u_j + K)}. \tag{A.2}$$

The argument in Section A.1 then yields $F_i / \sum_{j \in I} F_j = e^{-\alpha \delta u_i} / \sum_{j \in I} e^{-\alpha \delta u_j}$, where $\alpha \geqslant 0$.

A.3 Calculation of fixation probabilities

Here, we briefly outline the Markov chains defined by the two examples considered in the text, and how they are used to calculate fixation probabilities. Each of these chains is defined on some state space, S, and is specified by probabilities $P_{\mathbf{x} \to \mathbf{y}}$ for $\mathbf{x}, \mathbf{y} \in S$, representing the probability of transitioning from $\mathbf{x}$ to $\mathbf{y}$. Let $\mathbf{C}$ and $\mathbf{D}$ denote the monomorphic all-producer and all-non-producer states, respectively. The probability of reaching $\mathbf{C}$ from state $\mathbf{x}$ may be found by numerically solving the recurrence relation $\rho_{\mathbf{x}} = \sum_{\mathbf{y}} P_{\mathbf{x} \to \mathbf{y}} \rho_{\mathbf{y}}$, together with the boundary conditions $\rho_{\mathbf{C}} = 1$ and $\rho_{\mathbf{D}} = 0$. The probability of reaching the complementary state, $\mathbf{D}$, is then just $1 - \rho_{\mathbf{x}}$.

A.3.1 Cycle with *dB* updating

Since the cycle is one-dimensional and invading mutants spread in a cluster, the state space for the process may be reduced to $\{0, 1, \ldots, N\}$, which has size $N + 1$. In particular, a state α represents the number of C in a cluster, with the remaining $N - \alpha$ individuals forming a cluster of D individuals. Provided α is neither 0 nor N (the absorbing states), the state can transition from α to $\alpha - 1$ or $\alpha + 1$; otherwise, it remains in α. Since we allow for stochastic payoffs, we have to

keep track of the fecundity of each individual. Let $F_i(\alpha)$ denote the fecundity of individual i in state α, which is allowed to be a random variable and correlated with $F_j(\alpha)$ for $j \neq i$. To calculate the probability of gaining a mutant in state α under *dB* updating, we first calculate the probability that such a transition happens for a fixed fecundity vector $(F_i(\alpha))_{i=1}^N$. To do so, we only need to examine the boundaries between C and D on the cycle. Provided $0 < \alpha < N$, there are two such boundaries, although the individuals in these two boundaries overlap when $\alpha = 1$ or $\alpha = N - 1$.

As an example, suppose that $1 < \alpha < N - 1$. On one of the two boundaries, suppose that $i - 1$ is a C and i is a D. To gain a C along this boundary, increasing the state from α to $\alpha + 1$, i must first be selected for death, which happens with probability $1/N$. Then, $i - 1$ (type C) must win in a competition for reproduction against $i + 1$ (type D). The latter happens with probability $F_{i-1}(\alpha) \,/\, (F_{i-1}(\alpha) + F_{i+1}(\alpha))$. The other boundary, as well as the cases with $\alpha = 1$ or $\alpha = N - 1$, are analyzed similarly–except, in the latter case, a single individual is part of both boundaries. Finally, to get the overall transition probabilities, we take expectations over the transition probabilities that were calculated for fixed $(F_i(\alpha))_{i=1}^N$, where the source of randomness is the stochasticity in the payoffs that lead to fecundity.

A.3.2 Star with *Bd* updating

Calculations on the star are done similarly, except that in this case the state space is $\{0, 1\} \times \{0, 1, \ldots, N - 1\}$, which has size $2N$. An element (α, β) in this space represents whether the individual at the hub is a producer ($\alpha = 1$) or a non-producer ($\alpha = 0$), as well as how many individuals of type C are among the leaves, β. Under *Bd* updating, the probability that i reproduces in state (α, β) is $F_i(\alpha, \beta) \,/\, \sum_{j=1}^N F_j(\alpha, \beta)$, and the offspring replaces a randomly-chosen neighbor. As before, we are interested in the expected value of this probability, taken with respect to the distribution over $(F_i(\alpha))_{i=1}^N$ due to payoff stochasticity.

The code used to perform these calculations, which builds each of the transition matrices explicitly, is available at https://github.com/alexmcavoy/intermediate-selection.

References

[1] K. Sigmund. *The calculus of selfishness*. Princeton University Press, 2010.

[2] R. L. Trivers. The Evolution of Reciprocal Altruism. *The Quarterly Review of Biology*, 46 (1):35–57, 1971. doi: 10.1086/406755.

[3] J. Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982. doi: 10.1017/cbo9780511806292.

[4] R. Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.

[5] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998. doi: 10.1017/cbo9781139173179.

[6] G. Wild and A. Traulsen. The different limits of weak selection and the evolutionary dynamics of finite populations. *Journal of Theoretical Biology*, 247(2):382–390, 2007. doi: 10.1016/j.jtbi.2007.03.015.

[7] W. D. Hamilton. The genetical evolution of social behaviour, I. *Journal of Theoretical Biology*, 7(1):1–16, Jul 1964. doi: 10.1016/0022-5193(64)90038-4.

[8] W. D. Hamilton. The genetical evolution of social behaviour, II. *Journal of Theoretical Biology*, 7(1):17–52, Jul 1964. doi: 10.1016/0022-5193(64)90039-6.

[9] H. Ohtsuki, C. Hauert, E. Lieberman, and M. A. Nowak. A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441(7092):502–505, 2006. doi: 10.1038/nature04605.

[10] P. D. Taylor, T. Day, and G. Wild. Evolution of cooperation in a finite homogeneous graph. *Nature*, 447(7143):469–472, 2007. doi: 10.1038/nature05784.

[11] Y.-T. Chen. Sharp benefit-to-cost rules for the evolution of cooperation on regular graphs. *The Annals of Applied Probability*, 23(2):637–664, 2013. doi: 10.1214/12-aap849.

[12] H. Ohtsuki, J. M. Pacheco, and M. A. Nowak. Evolutionary graph theory: Breaking the symmetry between interaction and replacement. *Journal of Theoretical Biology*, 246(4):681–694, 2007. doi: 10.1016/j.jtbi.2007.01.024.

[13] B. Allen, G. Lippner, Y.-T. Chen, B. Fotouhi, N. Momeni, S.-T. Yau, and M. A. Nowak. Evolutionary dynamics on any population structure. *Nature*, 544(7649):227–230, 2017. doi: 10.1038/nature21723.

[14] E. Lieberman, C. Hauert, and M. A. Nowak. Evolutionary dynamics on graphs. *Nature*, 433 (7023):312–316, 2005. doi: 10.1038/nature03204.

[15] P. A. P. Moran. Random processes in genetics. *Mathematical Proceedings of the Cambridge Philosophical Society*, 54(01):60, 1958. doi: 10.1017/s0305004100033193.

[16] M. A. Nowak, A. Sasaki, C. Taylor, and D. Fudenberg. Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428(6983):646–650, 2004. doi: 10.1038/nature02414.

[17] H. Ohtsuki and M. A. Nowak. Evolutionary games on cycles. *Proceedings of the Royal Society B: Biological Sciences*, 273(1598):2249–2256, 2006. doi: 10.1098/rspb.2006.3576.

[18] H. Ohtsuki and M. A. Nowak. The replicator equation on graphs. *Journal of Theoretical Biology*, 243(1):86–97, 2006. doi: 10.1016/j.jtbi.2006.06.004.

[19] Y.-T. Chen, A. McAvoy, and M. A. Nowak. Fixation Probabilities for Any Configuration of Two Strategies on Regular Graphs. *Scientific Reports*, 6(1), 2016. doi: 10.1038/srep39181.

[20] D. Cassese. Replicator equation on networks with degree regular communities. *Applied Network Science*, 3(1), 2018. doi: 10.1007/s41109-018-0083-2.

[21] B. Allen, G. Lippner, and M. A. Nowak. Evolutionary games on isothermal graphs. *Nature Communications*, 10(1), 2019. doi: 10.1038/s41467-019-13006-7.

[22] C. Hauert and J. Miękisz. Effects of sampling interaction partners and competitors in evolutionary games. *Physical Review E*, 98(5), 2018. doi: 10.1103/physreve.98.052301.

[23] F. Débarre, C. Hauert, and M. Doebeli. Social evolution in structured populations. *Nature Communications*, 5, 2014. doi: 10.1038/ncomms4409.

[24] H. B. D. Kettlewell. Selection experiments on industrial melanism in the lepidoptera. *Heredity*, 9(3):323–342, 1955. doi: 10.1038/hdy.1955.36.

[25] H. B. D. Kettlewell. Further selection experiments on industrial melanism in the lepidoptera. *Heredity*, 10(3):287–301, 1956. doi: 10.1038/hdy.1956.28.

[26] M. E. N. Majerus. *Melanism: Evolution in Action*. Oxford University Press, Oxford, UK, 1998.

[27] M. E. N. Majerus. Industrial Melanism in the Peppered Moth, *Biston betularia*: An Excellent Teaching Example of Darwinian Evolution in Action. *Evolution: Education and Outreach*, 2 (1):63–74, 2008. doi: 10.1007/s12052-008-0107-y.

[28] L. M. Cook and I. J. Saccheri. The peppered moth and industrial melanism: evolution of a natural selection case study. *Heredity*, 110(3):207–212, 2012. doi: 10.1038/hdy.2012.92.

[29] A. S. Jump, J. M. Hunt, J. A. Martínez-Izquierdo, and J. Peñuelas. Natural selection and climate change: temperature-linked spatial and temporal trends in gene frequency in Fagus sylvatica. *Molecular Ecology*, 15(11):3469–3480, 2006. doi: 10.1111/j.1365-294x.2006.03027.x.

[30] A. Husby, M. E. Visser, and L. E. B. Kruuk. Speeding Up Microevolution: The Effects of Increasing Temperature on Selection and Genetic Variance in a Wild Bird Population. *PLoS Biology*, 9(2):e1000585, 2011. doi: 10.1371/journal.pbio.1000585.

[31] P. M. Altrock, A. Traulsen, and M. A. Nowak. Evolutionary games on cycles with strong selection. *Physical Review E*, 95(2), 2017. doi: 10.1103/physreve.95.022407.

[32] M. van Veelen and M. A. Nowak. Multi-player games on the cycle. *Journal of Theoretical Biology*, 292:116–128, 2012. doi: 10.1016/j.jtbi.2011.08.031.

[33] B. Skyrms. *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press, Dec 2003. doi: 10.1017/cbo9781139165228.

[34] R. Sugden. *The Economics of Rights, Co-operation, and Welfare*. B. Blackwell, 1986.

[35] W. Maciejewski, F. Fu, and C. Hauert. Evolutionary Game Dynamics in Populations with Heterogenous Structures. *PLoS Computational Biology*, 10(4):e1003567, 2014. doi: 10.1371/journal.pcbi.1003567.

[36] M. Nowak and K. Sigmund. The evolution of stochastic strategies in the Prisoner's Dilemma. *Acta Applicandae Mathematicae*, 20(3):247–265, 1990. doi: 10.1007/bf00049570.

[37] H. Ohtsuki. Evolutionary games in Wright's island model: kin selection meets evolutionary game theory. *Evolution*, 64(12):3344–3353, 2010. doi: 10.1111/j.1558-5646.2010.01117.x.

[38] H. Ohtsuki. Evolutionary dynamics of n-player games played by relatives. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1642):20130359–20130359, 2014. doi: 10.1098/rstb.2013.0359.

[39] A. McAvoy, B. Allen, and M. A. Nowak. Social goods dilemmas in heterogeneous societies. *Nature Human Behaviour*, 4(8):819–831, 2020. doi: 10.1038/s41562-020-0881-2.

[40] M. Doebeli, Y. Ispolatov, and B. Simon. Towards a mechanistic foundation of evolutionary theory. *eLife*, 6, 2017. doi: 10.7554/elife.23804.

[41] A. Traulsen, N. Shoresh, and M. A. Nowak. Analytical Results for Individual and Group Selection of Any Intensity. *Bulletin of Mathematical Biology*, 70(5):1410–1424, 2008. doi: 10.1007/s11538-008-9305-6.

[42] C. S. Gokhale and A. Traulsen. Evolutionary games in the multiverse. *Proceedings of the National Academy of Sciences*, 107(12):5500–5504, 2010. doi: 10.1073/pnas.0912214107.

[43] B. Wu, P. M. Altrock, L. Wang, and A. Traulsen. Universality of weak selection. *Physical Review E*, 82(4), 2010. doi: 10.1103/physreve.82.046106.

[44] B. Wu, J. García, C. Hauert, and A. Traulsen. Extrapolating Weak Selection in Evolutionary Games. *PLoS Computational Biology*, 9(12):e1003381, 2013. doi: 10.1371/journal.pcbi.1003381.

[45] D. Cooney, B. Allen, and C. Veller. Assortment and the evolution of cooperation in a Moran process with exponential fitness. *Journal of Theoretical Biology*, 409:38–46, 2016. doi: 10.1016/j.jtbi.2016.08.026.

[46] R. Durrett. Spatial evolutionary games with small selection coefficients. *Electronic Journal of Probability*, 19(0), 2014. doi: 10.1214/ejp.v19-3621.

[47] G. Szabó and C. Tőke. Evolutionary prisoner's dilemma game on a square lattice. *Physical Review E*, 58(1):69–73, 1998. doi: 10.1103/physreve.58.69.

[48] B. Wu, B. Bauer, T. Galla, and A. Traulsen. Fitness-based models and pairwise comparison models of evolutionary games are typically different—even in unstructured populations. *New Journal of Physics*, 17(2):023043, 2015. doi: 10.1088/1367-2630/17/2/023043.

[49] M. A. Nowak. *Evolutionary Dynamics: Exploring the Equations of Life*. Belknap Press, 2006.

[50] C. E. Tarnita and P. D. Taylor. Measures of Relative Fitness of Social Behaviors in Finite Structured Population Models. *The American Naturalist*, 184(4):477–488, 2014. doi: 10.1086/677924.

[51] B. Allen, C. Sample, P. Steinhagen, J. Shapiro, M. King, T. Hedspeth, and M. Goncalves. Fixation probabilities in graph-structured populations under weak selection. *PLOS Computational Biology*, 17(2):e1008695, 2021. doi: 10.1371/journal.pcbi.1008695.

[52] H. J. Muller. Artificial transmutation of the gene. *Science*, 66(1699):84–87, 1927. doi: 10.1126/science.66.1699.84.

[53] D. Atri and A. L. Melott. Biological implications of high-energy cosmic ray induced muon flux in the extragalactic shock model. *Geophysical Research Letters*, 38(19), 2011. doi: 10.1029/2011gl049027.

[54] N. Globus and R. D. Blandford. The Chiral Puzzle of Life. *The Astrophysical Journal*, 895 (1):L11, 2020. doi: 10.3847/2041-8213/ab8dc6.

[55] J. Tkadlec, A. Pavlogiannis, K. Chatterjee, and M. A. Nowak. Population structure determines the tradeoff between fixation probability and fixation time. *Communications Biology*, 2, 2019. doi: 10.1038/s42003-019-0373-y.

[56] D. Fudenberg and L. A. Imhof. Imitation processes with small mutations. *Journal of Economic Theory*, 131(1):251–262, 2006. doi: 10.1016/j.jet.2005.04.006.

[57] B. Allen and A. McAvoy. A mathematical formalism for natural selection with arbitrary spatial and genetic structure. *Journal of Mathematical Biology*, 78(4):1147–1210, 2019. doi: 10.1007/s00285-018-1305-z.

[58] S. Lessard. Long-term stability from fixation probabilities in finite populations: New perspectives for ESS theory. *Theoretical Population Biology*, 68(1):19–27, 2005. doi: 10.1016/j.tpb.2005.04.001.

[59] O. D. King and J. Masel. The evolution of bet-hedging adaptations to rare scenarios. *Theoretical Population Biology*, 72(4):560–575, 2007. doi: 10.1016/j.tpb.2007.08.006.

[60] A. McAvoy and B. Allen. Fixation probabilities in evolutionary dynamics under weak selection. *Journal of Mathematical Biology*, 82(3), 2021. doi: 10.1007/s00285-021-01568-4.

[61] B. Allen and C. E. Tarnita. Measures of success in a class of evolutionary models with fixed population size and structure. *Journal of Mathematical Biology*, 68(1-2):109–143, 2014. doi: 10.1007/s00285-012-0622-x.

[62] M. Assaf and M. Mobilia. Metastability and Anomalous Fixation in Evolutionary Games on Scale-Free Networks. *Physical Review Letters*, 109(18), 2012. doi: 10.1103/physrevlett.109.188701.

[63] F. C. Santos, M. D. Santos, and J. M. Pacheco. Social diversity promotes the emergence of cooperation in public goods games. *Nature*, 454(7201):213–216, Jul 2008. doi: 10.1038/nature06940.

[64] R. Ibsen-Jensen, K. Chatterjee, and M. A. Nowak. Computational complexity of ecological and evolutionary spatial dynamics. *Proceedings of the National Academy of Sciences*, page 201511366, 2015. doi: 10.1073/pnas.1511366112.

[65] W. Rudin. *Principles of Mathematical Analysis*. International Series in Pure and Applied Mathematics. McGraw-Hill, 1976. ISBN 9780070856134.

[66] I. J. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, Cambridge, MA, USA, 2016. `http://www.deeplearningbook.org`.

[67] R. Kindermann and J. L. Snell. *Markov Random Fields and Their Applications*. American Mathematical Society, 1980. doi: 10.1090/conm/001.