

Genetic Models in Evolutionary Game Theory: The Evolution of Altruism

Hannah Rubin

Received: 1 April 2014 / Accepted: 18 December 2014 / Published online: 1 January 2015
© Springer Science+Business Media Dordrecht 2015

Abstract While prior models of the evolution of altruism have assumed that organisms reproduce asexually, this paper presents a model of the evolution of altruism for sexually reproducing organisms using Hardy–Weinberg dynamics. In this model, the presence of reciprocal altruists allows the population to evolve to a stable polymorphic population where the majority of organisms are altruistic. Further, adding stochasticity leads to even larger numbers of altruists, while adding stochasticity to an analogous asexual model leads to more selfish organisms. The contrast between these outcomes demonstrates why it may be important to pay attention to the underlying genetics of a population.

1 Genetic Models and the Evolution of Altruism

The use of Hardy–Weinberg dynamics, or genetic models, in evolutionary game theory has been relatively rare, but not without reason. In many circumstances a model of asexual reproduction can approximate the evolutionary trajectories of a sexually reproducing population quite nicely—the replicator dynamics and Hardy–Weinberg dynamics yield qualitatively similar results, and the replicator dynamics are simpler (Maynard Smith 1982). However, there are certain interesting cases where the two dynamics differ—in particular, when the heterozygote strategy has the highest average fitness¹ or in a game with a

¹ For example, Nash bargaining games are discussed in Tennant (1999).

Table 1 Payoffs for a prisoner's dilemma

	C	D
C	3, 3	-1, 4
D	4, -1	0, 0

C stands for cooperate, which is equivalent to performing an altruistic act, *D* stand for defect, which is equivalent to not performing an altruistic act

mixed Nash equilibrium.^{2,3} In these cases, genetic models can often be used to re-evaluate the analysis of the simpler asexual models, with importantly different results. In using game theoretic tools to explain biological evolution and its impact on our moral or cultural norms, we should pay attention to the underlying biological mechanisms responsible for changes in phenotype over time as modeling these mechanisms differently can lead to different evolutionary outcomes (in models). In particular, this paper will evaluate the role of reciprocal altruism in genetic models of the evolution of altruism and compare this to the role of reciprocal altruism in similar asexual models.

In Sect. 2, I provide a basic deterministic genetic model of an infinite population with random interactions and no mutation, demonstrating that a population with a large number of altruists can be stable even in absence of a correlating mechanism. This is followed by an extension of the model which accounts for stochastic effects, such as finite population and mutation, in Sect. 3. Then, in Sect. 4, the contrast between a deterministic genetic model and a deterministic asexual model using replicator dynamics will be provided, followed by a comparison of the genetic stochastic model with a stochastic asexual model. In Sect. 5, I will conclude.

2 Deterministic Genetic Model

An act of altruism, in the biological sense, is something that lowers the fitness of the actor and increases the fitness of another organism. For instance, suppose each organism may perform an action that decreases its own fitness by 1 and increases another's fitness by 4. The resulting payoff matrix, which summarizes the effects on fitness from only this one interaction, is shown in Table 1 and represents a prisoner's dilemma, the game typically employed in models exploring the evolution of altruism.⁴

The evolution of altruism⁵ is seemingly puzzling for evolutionary theory: how could a trait which decreases the fitness of an organism ever evolve given that natural selection chooses traits with high fitness? In these models, some sort of

² The mixed Nash equilibrium may not be reached because the frequencies of strategies required for the population to play this strategy on average may not be achievable Hardy–Weinberg proportions (the proportions of traits that are consistent with the Hardy–Weinberg principle, described below).

³ For example, Hawk–Dove is analyzed in Maynard Smith (1982). In simple asexual models, populations playing Hawk–Dove will evolve to a mixed Nash equilibrium.

⁴ These payoffs and their interpretation are taken from Sober and Wilson (1998), 81–82.

⁵ I am referring to specifically biological altruism rather than psychological altruism.

correlation between like-types is usually used to explain how altruism could evolve. When there is some mechanism ensuring that altruists interact more frequently with other altruists, altruism can take over a population. Kin selection, group selection, network interactions, and other mechanisms which promote correlation have been employed in explaining the evolution of altruism (Skyrms 1994; Sober and Wilson 1998; Nowak 2006a). In many cases these mechanisms allow a population composed entirely of altruists to be stable, otherwise they can lead to a stable polymorphic population—a population with more than one phenotype present—in which altruists exist.

The evolution of altruism has also often been explored in the context of an *iterated* prisoner's dilemma. When the game is iterated, there is a greater variety of strategies available. In particular, an organism can play a strategy that conditions on its counterpart's action in the previous iteration. This allows the model to include what are called reciprocal altruists, who cooperate with altruists, but punish defectors by refusing to cooperate at some point. The presence of reciprocal altruists can lead to high levels of cooperation. Although reciprocal altruism is correlational in a sense (if you look at each iteration of the repeated game independently, an individual who cooperates is more likely to play against another individual who cooperates), it nevertheless is not a correlating mechanism in the same sense as in the models mentioned above. That is, it does not change the interaction frequencies of the individuals playing the game. Yet explanations employing reciprocal altruism all face a similar problem—while altruism is possible, it is fragile. Once a population is taken over by reciprocal altruists it can then be taken over by (unconditional) altruists (Nowak 2006b; Young and Foster 1991).⁶ This in turn opens the door for selfishness as the resulting altruistic population cannot resist invasion by defectors. We will see how this situation can be prevented in genetic models.

I will now lay out a model where the game is iterated 10 times⁷ and the reciprocal altruists play tit-for-tat (TFT). An organism with this strategy cooperates in the first iteration of the game and then copies whatever move its counterpart made for the rest of the iterations. The three strategies mentioned above can be interpreted biologically as arising from different genotypes.⁸ In a model with diploid organisms, there are more possible gene combinations than in models with haploid organisms.⁹ In the model here, there will be two versions (alleles) of the gene that

⁶ This could be achieved by drift or by positing a small fitness advantage to altruists because they do not have to expend energy “remembering” the past iteration of interaction.

⁷ This assumption is not essential. Similar results can be achieved with as few as two iterations. The results will also be produced if we assume that there is some probability of the game ending after each iteration (since the population is infinite, the payoffs are just multiplied by the expected number of rounds to give an expected payoff), as long as the probability of each next round is sufficiently high.

⁸ A genotype is a description of genetic makeup regarding a particular gene of interest.

⁹ Diploid organisms have two copies of each gene because their chromosomes come in pairs, while haploid organisms only have one copy of each gene.

Table 2 Payoffs for prisoner's dilemma iterated ten times, with three strategies and a background fitness of 20

	A	RA	S
A	50, 50	50, 50	10, 60
RA	50, 50	50, 50	19, 24
S	60, 10	24, 19	20, 20

A stands for altruism, RA for reciprocal altruism, S for selfishness

influences behavior in prisoner's dilemma type scenarios: an allele for altruism (call this allele A_1) and an allele for selfishness (call this allele A_2). An individual with two altruistic alleles (denoted A_1A_1) will always cooperate and an individual with two selfish alleles (A_2A_2) will always defect. Individuals with two copies of the same allele are called homozygotes. In cases of intermediate inheritance, the heterozygote (one with both an altruistic and a selfish allele, A_1A_2) has an intermediate trait, something 'between' altruism and selfishness.¹⁰ In the current model, this intermediate trait will be TFT.^{11,12}

Table 2 shows the payoff table for the game considered here. The payoffs are drawn from 10 iterations of the game in Table 1. Because the payoffs are interpreted as 'fitness' in the model here, I adjust values by adding 20 to each entry to avoid negative fitness values.¹³ This is because an organism's fitness is often thought of as the number of offspring it is expected to have and it does not make sense to think of an organism having 'negative' offspring. Additionally, the strategies are re-labeled to reflect the above interpretations in terms of selfishness versus altruism.

I will now consider how a population playing this game would evolve. I will first look at the deterministic case, where the population is assumed to be infinite and the end states of the population can be calculated from any possible initial population composition. I will then extend the model to account for stochastic effects—a finite population will play the game with a small chance of mutation.

The expected payoffs (or fitnesses) for organisms in an infinite population can be given in the standard evolutionary game theoretic way, where A is the proportion of altruists, RA the proportion of reciprocal altruists, S the proportion of selfish organisms:

$$\text{Expected Payoff for A} = 50A + 50RA + 10S$$

$$\text{Expected Payoff for RA} = 50A + 50RA + 19S$$

$$\text{Expected Payoff for S} = 60A + 24RA + 20S$$

¹⁰ A famous example of intermediate inheritance is in primroses, where one homozygote's flowers are red, the other's are white, and the heterozygote's flowers are pink.

¹¹ With these three strategies, TFT is a Nash equilibrium of the game. If all possible strategies are allowed, this would not be the case. For further discussion, see Axelrod (2006).

¹² For discussion of why TFT was chosen as the intermediate trait, see footnote 21 and Section 5.

¹³ This can be thought of as "background fitness" that each organism has from the other activities in its life.

Each organism's payoff depends on the composition of the population and the payoffs received from the game.

However, in infinite populations of diploid organisms, the proportion of individuals with a given trait must be expressed in terms of the allele frequencies. According to the Hardy–Weinberg principle, $\text{freq}(A_1A_1) = p^2$, $\text{freq}(A_1A_2) = 2pq$, and $\text{freq}(A_2A_2) = q^2$, where p is the frequency of altruistic alleles and q the frequency of non-altruistic alleles in the population. The expected payoffs can be re-expressed as follows, now including the biological shorthand for fitness, w , with the subscripts 11 for A_1A_1 , 12 for A_1A_2 , and 22 for A_2A_2 organisms:

$$w_{11} = 50p^2 + 100pq + 10q^2$$

$$w_{12} = 50p^2 + 100pq + 19q^2$$

$$w_{22} = 60p^2 + 48pq + 20q^2$$

These equations show how the population composition, and thus the payoff for any given strategy, is constrained by the combinations of genes. That is, because genes recombine probabilistically, certain proportions of genotypes (and thus proportions of strategies) are no longer possible unless artificially established. For instance, if one started with a population composed entirely of reciprocal altruists, the next generation would automatically have a different composition—25 % altruists, 50 % reciprocal altruists, and 25 % selfish.

In this model I employ Hardy–Weinberg dynamics, rather than the standard replicator dynamics. This is because the model deals with sexually reproducing diploid organisms rather than haploid organisms reproducing by cloning (as is commonly assumed for simplicity in game theoretic models). Hardy–Weinberg dynamics take into account that a gamete contains only one allele and that a parent passes on either of its alleles with equal probability.¹⁴ For example, while a gamete from an A_1A_1 organism will contain the A_1 allele 100 % of the time, a gamete from an A_1A_2 organism will contain the A_1 allele 50 % of the time and the A_2 allele 50 % of the time. This also means that the offspring do not necessarily have the same genetic make-up and will not necessarily play the same strategy as their parents. Determining the likelihood that an offspring will have a certain genotype is a matter of multiplying the likelihoods of its receiving those alleles. For instance, if an A_1A_1 and an A_1A_2 were to mate, their offspring would be equally likely to be A_1A_1 ($1.0 \times 0.5 = 0.5$) or A_1A_2 ($1.0 \times 0.5 = 0.5$).

The Hardy–Weinberg dynamics thus track the frequencies of each gene in the population, rather than the frequencies of traits, in each new generation.¹⁵ This is captured by the following equations:¹⁶

¹⁴ I assume there are no factors influencing this probability, such as meiotic drive.

¹⁵ Some may question why this discrete time dynamics are used rather than continuous time dynamics. For reasons to consider discrete time rather than continuous time, see Cressman (1992).

¹⁶ Since $q = 1 - p$ this dynamic is really one-dimensional, but I will continue to use q in equations for clarity.

$$p_{\text{next}} = \frac{p(w_{11}p + w_{12}q)}{\bar{w}}$$

$$q_{\text{next}} = \frac{q(w_{22}q + w_{12}p)}{\bar{w}}$$

where \bar{w} , the average fitness of the population, is calculated in the intuitive way:

$$\bar{w} = w_{11}p^2 + w_{12}2pq + w_{22}q^2$$

Then, to find the rest points, we set $p = p_{\text{next}}$ and solve for p in order to find when the current population composition is the same as the next, or when selection pressures have stopped causing the population to evolve. This yields:

$$p = \frac{p(w_{11}p + w_{12}q)}{\bar{w}}$$

$$1 = \frac{w_{11}p + w_{12}q}{w_{11}p^2 + w_{12}2pq + w_{22}q^2}$$

$$p = \frac{w_{22} - w_{12}}{w_{11} - 2w_{12} + w_{22}}$$

Then, the payoffs from the current model can be inserted into this equation to obtain:

$$p = \frac{60p^2 + 48pq + 20q^2 - (50p^2 + 100pq + 19q^2)}{50p^2 + 100pq + 10q^2 - 2(50p^2 + 100pq + 19q^2) + 60p^2 + 48pq + 20q^2}$$

$$54p^2 + 45p + 1 = 0$$

Solving this yields two rest points, and there are two additional rest points because the dynamics used here satisfy extinction:¹⁷

1. $p = 0$
2. $p = -\frac{1}{36}\sqrt{201} + \frac{5}{12} (\approx 0.0228)$
3. $p = \frac{1}{36}\sqrt{201} + \frac{5}{12} (\approx 0.8104)$
4. $p = 1$

(1) represents an entirely selfish population (there is no altruistic gene present), while (4) represents a population composed entirely of altruists. (2) and (3) represent polymorphic states—both genes and thus all three phenotypes are present. In (2), the population composition is almost entirely selfish: roughly 95.5 % of individuals are selfish, 4.5 % of individuals are reciprocal altruists, and <1 % are

¹⁷ Once an allele is no longer part of the population, it does not get reintroduced into the population because there is no mutation or other mechanism to introduce new alleles. An allele that is extinct stays extinct.

altruists. Alternatively, (3) represents a population that is about 65.7 % altruist, 30.7 % reciprocal altruist, and 3.6 % selfish. In this population, almost all interactions will be altruistic in character. That is, in the cases where altruists interact with either other altruists or reciprocal altruists and where two reciprocal altruists interact, no individual will ever defect. These interactions can be considered altruistic—organisms always perform the altruistic act—and will constitute roughly 92.9 % of the interactions between organisms in (3).

Rest point stability can be determined by the relative fitnesses of the alleles.¹⁸ When the fitness of the A_1 allele is greater than the fitness of the A_2 allele, there will be a higher proportion of the A_1 allele in the next generation. That is, when the A_1 allele has a higher fitness, $p_{\text{next}} > p$ and the population is evolving to be more altruistic. Below I calculate the conditions under which this is the case.

$$w_{A1} > w_{A2}$$

$$\frac{w_{11}p^2 + w_{12}p(1-p)}{p^2 + p(1-p)} > \frac{w_{22}(1-p)^2 + w_{12}p(1-p)}{(1-p)^2 + p(1-p)}$$

$$p > \frac{w_{22} - w_{12}}{w_{11} - 2w_{12} + w_{22}}$$

The proportion of the altruistic allele is increasing when $54p^2 + 45p + 1 > 0$, between the two polymorphic states (2) and (3).

Figure 1 shows that although (1) is a locally asymptotically stable (LAS)¹⁹ population of all defectors, it has a very small basin of attraction.²⁰ A population will only evolve to such a state if the initial population composition is within the tiny basin of attraction, so this outcome is very unlikely. On the other hand, (3) attracts almost all initial population states, so the outcome of evolution acting on an arbitrary initial population should be expected to lead to a LAS polymorphic state in which the vast majority of interactions are altruistic.²¹

¹⁸ Alternatively, the eigenvalues of the Jacobian matrix can be used to check which rest points are stable. For (1) and (3), they are both negative [-1 and -0.05 for (1) and -1 and about -0.025 for (3)] so these rest points are stable. For (2) and (4), only one eigenvalue is negative, so these rest points are unstable.

¹⁹ If the population state is perturbed by a small amount in any direction (e.g. if there is some mutation causing the population to move away from the rest point), evolution will lead back to the rest point.

²⁰ A state's basin of attraction is the set of possible initial population compositions that evolve toward the state.

²¹ Of course, either the altruistic or selfish organisms could also be modeled as heterozygous. If heterozygotes are altruists (with reciprocal altruist and selfish homozygotes) then there are two stable equilibria: one composed entirely of reciprocal altruists and one composed entirely of selfish organisms, with the basin of attraction for reciprocal altruists just slightly larger. On the other hand, if the model assumes heterozygotes are selfish organisms then there is only one stable equilibrium composed entirely of reciprocal altruists.

I focus on the case where heterozygotes are reciprocal altruists for two reasons. First, the case where the heterozygote is a reciprocal altruist is in some sense the worst case for altruism—both of the other possible assumptions for the heterozygous genotype lead to stable equilibria composed of reciprocal altruists, where no organism ever defects. Second, the case where heterozygotes are reciprocal altruists is also arguably the most interesting because it is the only case where there is a stable equilibrium in which there are unconditional altruists rather than just reciprocal altruists.



Fig. 1 The population states of the dynamics with the rest points labeled (*open circles* for unstable rest points and *blue circles* for stable rest points). *Arrows* represent selective pressures. An *arrow* pointing to the *right* represents population compositions in which $w_{A1} > w_{A2}$ since in this case the population will evolve toward $p = 1$ and an *arrow* pointing to the *left* represents the opposite case. (Color figure online)

3 Stochastic Genetic Model

Stochastic effects, or ‘noise,’ can be added into the genetic model from Sect. 2 by considering a finite population in which there is a possibility of mutation. When modeling a finite population, the expected fitness calculations are slightly different than when modeling an infinite population. First, we can no longer calculate the frequency of each type of organism based on the probabilistic recombination of genes. This is possible in an infinite population model, but not in a finite population model where probabilities merely correspond to expectations and thus do not necessarily correspond to the distribution of phenotypes in the next generation. Second, in a finite model we must account for the fact that an organism cannot interact with itself. For instance, if there are 100 altruists in a population, each altruistic organism has 99 other altruists with which it could possibly interact. The expected payoffs for the finite population are as follows:

$$w_{11} = 50 \left(\frac{A \cdot N - 1}{N} \right) + 50RA + 10S$$

$$w_{12} = 50A + 50 \left(\frac{RA \cdot N - 1}{N} \right) + 19S$$

$$w_{22} = 60A + 24RA + 20 \left(\frac{S \cdot N - 1}{N} \right)$$

where N is the population size and A , RA , and S now refer to the numbers of altruists, reciprocal altruists, and selfish organisms rather than the proportions of these organisms. The average fitness is again calculated in the intuitive way:

$$\bar{w} = w_{11}A + w_{12}RA + w_{22}S$$

To capture the probabilistic nature of death and reproduction in a finite population, the Moran Process can be adapted to account for sexual reproduction (Moran 1962). In each time period, two individuals are chosen to reproduce with probability proportional to their fitness:

$$P(i = \text{1st parent}) = \frac{w_i}{\bar{w}} \quad \text{for all } i \in N$$




Fig. 2 Long-run behavior as a function of population size. Shown are the fractions of data points where the population is composed of at least 90 % altruists, labeled (4), and where the population is in a region around the polymorphic equilibrium, labeled (3). The region around the polymorphic equilibrium consists of all the population states where the percent of each type is ± 5 % of their percentage in the polymorphic equilibrium (3). There is no data for the selfish equilibrium (1) because, except for when $N = 33$, the population never reached 90 % selfish organisms

$$P(j = 2\text{nd parent}) = \frac{w_j}{\bar{w}_{-i}} \quad \text{for all } j \in N \text{ s.t. } j \neq i$$

where \bar{w}_{-i} is the average fitness of the population without the first parent. These organisms pass on either one of their genes with equal likelihood.

Additionally, I assume the mutation rate is $\mu = 0.001$. Note that mutations in diploid organisms amount to a change in a single allele and the probability of a gene mutating is independent of the other gene undergoing a mutation. At the end of each time period, one individual is randomly chosen to die in order to keep the population size fixed at N . This process was run for 1,000,000 rounds with data points taken every 100 rounds, with an initial population consisting of the three strategies in equal proportions (Fig. 2).

While, in the short run, a finite population with these stochastic effects evolves much like a population under a deterministic dynamic, in the long run the population may end up spending more time near a population state that had a relatively small basin of attraction in the deterministic model. That is, adding noise into the model can cause the population to reach certain states more frequently than might be expected based on the deterministic model.

In fact, introducing stochastic effects into the genetic model shows that it is possible for a population to spend time near the entirely altruistic rest point (4). This is because near (3) altruists have high fitness on average, close to that of the reciprocal altruists—altruists interact frequently with reciprocal altruists or other altruists, so are not very often taken advantage of by selfish organisms—so it is not difficult for stochastic effects to lead the population to (4) just by chance. The amount of time the population spends there is fairly small, but it is significant

especially considering that, except for when $N = 33$, the population never gets sufficiently close to the all selfish equilibrium. The population also spends much of its time around the polymorphic equilibrium (3).

As noise decreases, the population spends more time around the polymorphic equilibrium, but this is unsurprising. This state's basin of attraction encompasses nearly all possible starting states in the deterministic model, so reducing noise just lets these selective pressures have more impact. More remarkably, the time spent near an all altruist population (4) increases as the noise decreases. With lower levels of noise, it is still possible to transition from (3) to (4) by chance, and once a population is near (4) it stays there much longer than it would with higher levels of noise. This is because near (4), there is hardly any fitness difference between altruists and reciprocal altruists, so the selective pressure leading towards (3) is very weak. Of course, transitioning from (3) to (4) is less likely than transitioning from (4) to (3) because there is selective pressure leading the population toward (3), which is why the population spends more time around (3) in the long run.

4 Reciprocal Altruism in Asexual Models

We are now in the position to investigate how a genetic model compares with other similar models of the evolution of altruism. Since Axelrod's famous experiments, where TFT emerged as the clear winner, many have investigated the advantage of TFT in the prisoner's dilemma (Axelrod 2006; Nowak 2006a). To highlight the differences attributable to the change in underlying dynamic, I will discuss in detail asexual models with the same three strategies as were used in the previous section: always cooperate (altruist), TFT (reciprocal altruist), and always defect (selfish). I will first look at the deterministic case where the behavior of the population can be computed directly from the starting conditions. I will then describe a stochastic model, in which chance events play a role in the evolutionary dynamic.

4.1 Deterministic Asexual Model

As mentioned, the replicator dynamics is commonly used for modeling asexual reproduction in an infinite population. In this dynamic the change in p , now interpreted as the frequency of a strategy i , is proportional to the difference between the strategy's fitness and the average fitness in the population:

$$\dot{p} = p[w_i - \bar{w}]$$

If the fitness of a strategy is greater than the average fitness of the population, the frequency of that strategy will increase.

In an asexual model with random encounters and no mutations in an infinite population, there is only one LAS outcome: a completely selfish population (S-equilibrium). A population of all reciprocal altruists (RA-equilibrium) is also stable, but in a less strict sense. The RA-equilibrium is *neutrally* stable—if the population is perturbed, it will remain nearby the RA-equilibrium. One might think that




Fig. 3 The simplex for the replicator dynamics in an infinite population with random interactions and no mutations. Payoffs are as in Table 2. The basin of attraction for selfish individuals is along the *upper right axis* and the basin of attraction for reciprocal altruists is in the *lower left of the triangle* (Sandholm 2012)

evolution would always lead back to the RA-equilibrium if the population state is moved slightly away from it. However, if the perturbation introduces altruists, the population can get stuck at one of an infinite number of neutrally stable polymorphic states consisting of both altruists and reciprocal altruists that exist outside the basin of attraction for selfishness (S-basin). This is because in the absence of defectors, both of these strategies have the same fitness, and thus there is no selective pressure to lead the population back to the RA-equilibrium (Fig. 3).

The differences between the two dynamics are immediately noticeable. With the replicator dynamics, there is a large basin of attraction for a population composed entirely of reciprocal altruists (RA-basin), as there are a number of initial population compositions that evolve toward the RA-equilibrium. In contrast, a population composed entirely of reciprocal altruists is not possible in the deterministic genetic model provided in Sect. 2. Additionally, the S-basin is substantial under the replicator dynamics, whereas in the genetic model only a tiny fraction of possible starting states evolve toward the selfish equilibrium (1).

Furthermore, under the replicator dynamics with these assumptions there is not a LAS equilibrium where there are any (unconditional) altruists. This is because any rest point in which altruists appear—those composed of altruists and reciprocal altruists inside the RA-basin—is only neutrally stable. However, under the Hardy–Weinberg dynamics the LAS polymorphic state (3) is dominated by altruists. It is worthwhile to note that this polymorphic state is not sustained by some underlying population structure (as is the case in models utilizing networks or graphs with the replicator dynamics). Because heterozygote superiority allows polymorphic states to be stable, sexual reproduction offers opportunity for maintaining genetic diversity that asexual reproduction does not.

4.2 Stochastic Asexual Model

Young and Foster (1991) incorporate stochastic effects into a population playing an iterated prisoner's dilemma with these three strategies. The model is of a finite population in which initially there are equal proportions of each strategy. Stochastic effects are introduced in two ways. First, there are a fixed number of encounters in each time period but instead of a fixed number of iterations in each encounter, there is a certain probability that the game will continue after each iteration. This means, for instance, that the number of times a selfish organism can take advantage of an altruist per encounter can vary widely and therefore so can its payoff. Second, mutations are introduced into the model so that the simplex is bounded away from the edges; a population can no longer evolve toward an equilibrium and get stuck there indefinitely. They assume a mutation rate of $\mu = 0.001$, as was assumed in the stochastic genetic model in Sect. 3. Unlike in a genetic model where mutations in diploid organisms amount to a change in a single allele, mutations in haploid organisms correspond to an entirely new genotype. Young and Foster show that in the long run the population can be expected to spend much of its time in the selfish equilibrium and that, surprisingly, the effect increases as noise decreases.

Young and Foster explain that the population ends up being near the S-equilibrium more often in the long run because escaping the RA-basin is more likely. Getting out of the RA-basin is easier than escaping S-basin for two reasons. First, the selection pressure for reciprocal altruists over altruists within the RA-basin is not as strong as the selection pressure for selfish individuals over altruists within the S-basin. An altruistic mutant arising close to the S-equilibrium is likely to be taken advantage of quickly and die out. On the other hand, an altruistic mutant is not selected against as strongly near the RA-equilibrium because there are not many selfish organisms. The altruistic mutant can be expected to survive longer, meaning there is a longer time window in which successive mutations can occur before the previous altruistic mutants die out. Second, during this time window, it is possible that altruist–altruist interactions are longer than interactions between reciprocal altruists. Altruists could benefit from the longer interactions enough to make up for the harm caused by an occasional interaction with a selfish organism, meaning altruists could actually have higher fitness than reciprocal altruists within the RA-basin. Defectors are still selected against in the RA-basin, so with some positive probability the population could evolve along the edge of the simplex until it crosses into the S-basin where the population would be quickly swept up into the S-equilibrium because the selective pressure for selfish organisms over altruistic organisms is so strong.

Let us compare these findings to the stochastic genetic model from Sect. 3. Interestingly, adding stochastic effects to the genetic model does not yield the result that the population spends more time near the entirely selfish population. In fact, adding stochastic effects into the genetic model causes the population to spend more time as an approximately *all altruist* population. For similar reasons as those in the asexual case, it is not too difficult to get to a population of all altruists from the

polymorphic equilibrium (3).²² This is because around (3), altruists and reciprocal altruists have roughly the same fitness (there are very few selfish individuals to prey on the altruists). However, a small number of defectors introduced by chance mutations in the all-altruist population will lead the population back towards (3), not to the selfish equilibrium (1). This is in contrast to the asexual model, in which a small number of defectors introduced into an all-altruist population will take over—the population will quickly evolve toward the S-equilibrium. In the stochastic genetic model the population the population never gets close to the S-equilibrium.

Also recall that the time spent near an all altruist population increases as the noise decreases. While these results may seem in stark contrast to Young and Foster's, the two findings are more similar in character than they at first appear. The common result is that, under a stochastic process, certain population states which might be unlikely in the short term can be far more likely in the long term. Further, the surprising effects of introducing a stochastic element are amplified as the stochastic element is diminished.

5 Conclusion

While there are often good reasons why the Hardy–Weinberg dynamics are passed over in favor of the simpler, more manageable replicator dynamic, we have seen that genetic models should not be neglected by game theorists entirely. The introduction of diploid organisms with heterozygous reciprocal altruists can lead to a population with high levels of cooperation without the need for correlation between types. Further, stochastic effects can actually lead to higher levels of cooperation in the long-run. These sorts of (perhaps unexpected) consequences of an added layer of realism in biological models suggest that ignoring the underlying genetics of a trait may lead us astray.

The analysis can work with other intermediate strategies. For instance, in this model TFT is equivalent to GRIM, which cooperates on the first iteration then defects in every iteration after their counterpart has defected. Analysis for strategies that start defecting but cooperate with cooperators in subsequent iterations yields analogous results. It seems that the only requirement for a heterozygote strategy that allows the analysis to go through is that it somehow conditions its behavior on its counterpart's strategy in such a way that it ends up cooperating with altruists and defecting against selfish individuals.

One might think that a more realistic assumption would be that the intermediates play a mixed strategy or that they simply alternate between cooperation and defection. However, if we think of the genes as encoding for a disposition towards certain types of action, rather than a specific strategy in a specific game (which is perhaps a more realistic assumption), then the heterozygote can be seen as having a disposition for both altruism and selfishness. We might say then that an organism

²² It's worth noting that for these simulations, there were a fixed number of iterations per engagement. It is easy to see why the conclusions drawn here would extend to those cases with further stochastic effects: the role the varying length of engagement played was in explaining how altruists could easily accumulate by chance along the edge of the RA-basin. Being drawn to a population with more altruists only makes in more likely to stay around the all altruist equilibrium. See Young and Foster (1991), 152.

using TFT is disposed be altruistic, but also to be selfish when the situation calls for it. One might also object that assuming the heterozygote plays TFT is unrealistic because many organisms might not have the cognitive capacities to implement such a strategy. However, there is biological evidence that a variety of organisms can play strategies as cognitively complex as TFT. For instance, birds have been found to have sufficient cognitive ability (Krama et al. 2012). There is even some evidence that still simpler organisms, such as small fish like sticklebacks and guppies, can carry out TFT-like strategies (Dugatkin 1991).

While it is plausible that an organism could play TFT without a high level of cognitive ability, the fact that the heterozygote, rather than a homozygote, is assumed to play TFT is not meant as an attempt to make the model more biologically realistic. Instead, this model is like many other biological models of the evolution of altruism which are highly idealized and meant to demonstrate possible outcomes of evolution given certain reasonable assumptions, rather than to describe real populations. Supposing TFT is played by the heterozygote leads to interesting results—a stable population dominated by altruists—so it clearly demonstrates why utilizing genetic models is of interest to evolutionary game theorists. Although assigning different strategies to the heterozygote may not always have as profound of an effect on evolutionary outcomes, the results presented in this paper point to a whole class of models which have been overlooked in many game-theoretic contexts and which would be profitable to study.

This is not to argue that simplicity is the only benefit of the replicator dynamics over Hardy–Weinberg dynamics. There are, of course, situations where genetic modeling is inappropriate. Since the replicator dynamics are formally equivalent to many models of cultural evolution, models attempting to explain the evolution of altruism, cooperation, or related social phenomena using the replicator dynamics can apply to both biological and social evolution. If the aim is to show how people within a society adapt their behavior to their peers, a genetic model cannot explain this sort of horizontal transmission.

Rather, in evolutionary game theory genetic models are useful for explaining underlying biological drives or dispositions to behave in a certain manner. This paper has demonstrated that, even in very simple games like the iterated prisoner's dilemma, biological evolution in diploid organisms may lead to drastically different results from cultural (or asexual) evolution. This indicates that attention needs to be paid to genetic models in evolutionary game theory.

Acknowledgments I would like to thank Simon Huttegger, Brian Skyrms, Justin Bruner, Cailin O'Connor audiences at GIRLS 2014 and two anonymous reviewers for their helpful comments. I would also like to thank Zac Ernst, Paul Weirich, André Ariew, and Philip Robbins for their help on a previous version of this paper.

References

- Axelrod, R. (2006). *The evolution of cooperation* (2nd ed., pp. 27–69). New York, NY: Basic Books.
- Cressman, R. (1992). *The stability concept of evolutionary game theory* (pp. 58–59). Berlin: Springer.
- Dugatkin, L. A. (1991). Dynamics of the tit-for-tat strategy during predator inspection in the guppy (*poecilia-reticula*). *Behavioral Ecology and Sociobiology*, 29(2), 127–132.

- Krama, T., Vrublevska, J., Freeberg, T. M., Kullberg, C., Rantalaand, M. J., & Krams, I. (2012). You mob my owl, I'll mob yours: Birds play tit-for-tat game. *Scientific Reports*, 2, 1–3. doi:[10.1038/srep00800](https://doi.org/10.1038/srep00800).
- Moran, P. A. P. (1962). *The statistical processes of evolutionary theory*. Oxford: Clarendon.
- Nowak, M. A. (2006a). 5 Rules for the evolution of cooperation. *Science*, 314, 1560–1563.
- Nowak, M. A. (2006b). Evolutionary dynamics of cooperation. In *Proceedings of the International Congress of Mathematicians*, Madrid (pp. 1523–1540).
- Peyton, Y. H., & Foster, D. (1991). Cooperation in the short and in the long run. *Games and Economic Behavior*, 3, 145–156.
- Sandholm, W. H., Dokumaci, E., & Franchetti, F. (2012). Dynamo: Diagrams for evolutionary game dynamics. <http://www.ssc.wisc.edu/~whs/dynamo>.
- Skyrms, B. (1994). Darwin meets the logic of decision: Correlation in evolutionary game theory. *Philosophy of Science*, 61(4), 503–528.
- Smith, J. M. (1982). *Evolution and the theory of games* (pp. 41–43). Cambridge, UK: Press Syndicate of the University of Cambridge.
- Sober, E., & Wilson, D. S. (1998). *Unto others: The evolution and psychology of unselfish behavior* (pp. 1–194). Cambridge: Harvard University Press.
- Tennant, N. (1999). Sex and the evolution of fair dealing. *Philosophy of Science*, 66(3), 391–414.