In Love and War;
altruism, norm formation, and two different types of group selection

Matthijs van Veelen (corresponding author)

Department of Economis and Econometrics, Universiteit van Amsterdam

Roetersstraat 11, 1018 WB Amsterdam, the Netherlands

tel: ++31 20 525 5293 / fax: ++31 20 525 5283

C.M.vanVeelen@uva.nl

Astrid Hopfensitz

CISA - Swiss Center for Affective Sciences, University of Geneva

7 rue des Battoirs, 1205 Geneva, Switzerland

tel: ++41 22 37 99 808 / fax: ++41 22 379 98 44

Astrid.Hopfensitz@cisa.unige.ch

September 5, 2007

Abstract

We analyse simulations reported in "The co-evolution of individual behaviors and social institutions" by Bowles, Choi & Hopfensitz (2003) in the *Journal of Theoretical Biology* **223**, 135-147 and begin with distinguishing two types of group selection models. The literature does not provide different names for them, but they are shown to be fundamentally different and have quite different empirical implications. The working of the first one depends on the answer to the question '*is the probability that you also are an altruist large enough*', while the other needs an affirmative answer to '*are our interests enough in line*'. The first one therefore can also be understood as a kin selection model, while the working of the second can also be described in terms of direct benefits. The actual simulation model is a combination of the two. It is also a Markov chain, which has important implications for how the output data should be handled.

1 Introduction

The simulations by Bowles, Choi & Hopfensitz (2003) of the co-evolution of norms and altruistic behaviour are interesting for a number of reasons. One of them is that the social behaviour of the individuals that feature in the simulations is more complex and realistic than models normally allow for (see Lehmann & Keller, 2006, for an overview). Agents here may not only be altruistic towards members of their own group, but they also engage in wars with other groups. Furthermore there are norms that co-evolve with altruistic behaviour. Simulations that do encompass intergroup conflicts and evolution of norms point at reasons why human other-regarding behaviour as well as our norms have their peculiarities, one of which is that the altruism we display is rather situation specific. Experimental evidence supports the idea that humans behave much less altruistic towards non-group members than they do towards group members, even if individuals are randomly assigned to groups (Tajfel & Turner, 1986). Furthermore, interactions between groups are observed to be more competitive than between individuals and this is known as the interindividual-intergroup discontinuity effect (see for instance Wildschut et al, 2003). Performing simulations which explicitly model inter-group conflict therefore is a valuable thing to do. It may not only clarify why people evolved to be altruistic, but also why altruism is not lavished on everyone and in all circumstances.

The actual form of the simulation model will shortly be discussed in Section 2, but it is worthwhile mentioning here that it is not set up as a special case, numerical version or interesting variant of an existing model. The decisive argument for the choice of a simulation model seems to have been whether or not it is realistic, as for instance the careful considerations concerning the choice of benchmark values of the parameters indicate. This makes it an interesting model, but it might also be the reason why it is not so easy to pinpoint the features of the simulation model that are essential for the results. The first aim of this paper is therefore to disentangle the different model ingredients. In Sections 3 and 4 we distinguish two models that both are candidates for the label "group selection model", but that are nonetheless very different. The first one could be called the standard group selection model (see Price (1972), Hamilton (1975), Frank (1998) and Sober & Wilson (1998)) and there altruism evolves as a result of assortative group formation. In this type of model, the degree to which group composition varies across groups determines for which combinations of costs and benefits altruistic behaviour is selected. The essential feature of the second type of model is that the fate of group members are, to some extent, aligned. It therefore does not depend on or need assortative group formation to work. In Section 4 we show how this aspect of the simulations can be described as an extension of a model by Weibull & Salomonsson (2006). We also point out that the first type of model fits within the kin selection framework, while the working of

the second can be described in terms of direct benefits.

The simulations may be a combination of the two, but it is very much worth exploring how the two composing models differ in their predictions of human behaviour. While the first type of group selection model for instance can be used to explain racial discrimination (Hamilton, 1975), the behaviour that emerges from the second type - and that is found in the experiments by Tajfel & Turner (1986) - is better described as groupishness (see Ridley, 1997). Another implication of the first type of model is that relatedness (genetic similarity) within groups must be high enough to explain altruistic behaviour for given costs and benefits. This is in contrast with the second model, which does not imply such a threshold. Section 5 discusses these empirical implications as well as some experimental evidence.

Then we put the two model ingredients together and return to the actual simulations with a rather technical section. This section characterizes the dynamics in the simulation model and in subsection 6.1 we show that the simulations are in fact Markov chains. Realizing this helps us properly interpret the numerical output such as for instance the average frequency of altruists over a simulation run. Then subsection 6.2 shows that an 'estimation of the Price equation' as done in BCH is not a meaningful exercise, because what it does is fit a misspecified econometric model to the simulation data.

Although parts of this article are written as critical comments on the analysis in BCH, this must emphatically not be taken as an argument against the importance of their simulation results. On the contrary, we hope that pointing out what drives the evolution of altruism and norms in the simulations is helpful in understanding why they are so interesting and thoroughly relevant.

2   The simulations

For a full description of the simulations, we refer to the original article. Here we only give a brief sketch.

In the simulations, there is a fixed number of groups and a fixed total number of individuals. Individuals can be altruistic (A) or not altruistic (N). Initially, individuals are randomly assigned to groups, implying that not all groups will contain the same number of individuals. Then every round consists of a number of stages. *1. Pairing.* First, group members are paired within their groups to play a PD game, where the A-players play cooperate and the N-players play defect. Cooperation means that a benefit b is conferred to the other player at cost c. The resulting payoffs are used in two further stages, namely 2. Reproduction and 5. Group competition. *2. Reproduction.* The next generation in each group is produced by drawing from a binomial distribution, with the probability of drawing

3

an A player equal to the total payoffs of A-players in the group in the old generation, divided by the total payoff of all players in the group. *3. Mutation.* Each new generation member is drawn with probability e to redetermine its type, where it will be A or N with equal probability. *4. Migration.* With probability m each member of the new generation relocates to a randomly selected other group. *5. Group competition.* With probability k each group is selected and the groups that are drawn are paired randomly. The winning group is the one with the highest total payoff, determined in the first stage. *6. Repopulation and fission.* The losing group members disappear. The winning group is temporarily enlarged to the total size of both groups, where the new group members have the same frequency of A-players as the winning group, with some roundoff error. The resulting group splits with members assigned randomly to the two new groups, that therefore can differ in size.

In the versions with norms, a group can be endowed with either a sharing norm or a segmentation norm, or both. These norms take values between 0 and 1, where the sharing norm is modelled as a linear tax t, collected from the members payoffs, with the proceeds distributed equally to all members of the population.[1] A segmentation norm ensures that members interact with a similar type with probability s and are paired randomly with probability $1 - s$. Norms in this setting therefore can be seen as anything that affects the expected payoffs of the different types by changing (the payoffs in) the game. After a group conflict, the new groups adopt the norms, that is, the values for s and t, of the winning group. The norms are introduced by small random changes in values of t and s.

The BCH paper contains a thorough description of the parameter space, indicating that small groups, low migration rates and high frequencies of group conflict are good for the evolution of altruism, and that the possibility of institutions enlarge the part of the parameter space that is favourable for altruism. The latter may seem obvious, once we see it; both norms decrease the (expected) payoff of N-players and increase the (expected) payoff of A players, thereby reducing selection pressure against A-players. Yet it is a very valuable insight that, even if the norms come at a cost, they themselves can spread too. The reason is that norms elevate the frequencies of A-players within the group (or, rather, they reduce the speed of the demise) and therefore they can help the groups that have them win inter-group conflicts, by means of which the norm itself also spreads.

---

[1] BCH state that this reduces payoff differences between A's and N's to $(1-t)c$, but a more accurate description is that within a group, it reduces the expected difference in payoffs between A's and N's to $(1-t)\left(c + \frac{b}{n}\right)$, where n is the size of the group.

3   The standard group selection model

Section 2 of BCH contains a simple model of multi-level selection that we will categorize as an example of the standard group selection model. We will shortly characterize it and mention a common misunderstanding concerning the workings of the model. In this model, the population is divided into groups. Within these groups, members can, at cost $c > 0$, choose to confer a benefit $b > c$ to a randomly chosen member of the group the individual is in. Here we are to think of b and c as differences in fitness and in this version altruism is a binary trait; an individual is either an A-player that does perform this (altruistic) act or an N-player that does not. It can be useful to let p denote the fraction of A-players in the whole population and p_j the fraction of A-players in group j.

Whether or not the fraction of A-players in the population will rise, crucially depends on how the total population is divided into groups. If this division is completely assortative, so that there are only homogeneous groups, then obviously the fraction of A-players in the population will rise; all N-players do not give and do not receive, and all A-players do give and do - in expectation - receive once, so they have an expected net gain of $b - c$. If on the other hand group formation is completely random, then the fraction of A-players will fall in expectation; all players receive b with probability p, but A-players lose c where N-players lose nothing. More generally, the expected payoff to an A-player is $b\,\mathbb{P}\left(A \mid A\right) - c$, where $\mathbb{P}\left(A \mid A\right)$ is the probability of being paired with an A-player if you are an A-player yourself. The expected payoff to an N-player is $b\,\mathbb{P}\left(A \mid N\right)$, where $\mathbb{P}\left(A \mid N\right)$ is the probability of being paired with an A-player if you are an N-player. The share of A-players rises if the first exceeds the second, that is, if $b\left\{\mathbb{P}\left(A \mid A\right) - \mathbb{P}\left(A \mid N\right)\right\} - c > 0$, where $\mathbb{P}\left(A \mid A\right) - \mathbb{P}\left(A \mid N\right)$ can be seen as a general measure of either relatedness or assortativeness of group formation. This formula obviously is Hamilton's rule. Following the classification of models by Lehmann and Keller (2006), it is clear that any model that implies that this difference in probabilities diverges from zero falls in the category of kin selection (see also Lehmann et al., 2007 and Grafen, 2007. Both these papers provide a parallel analysis in terms of kin selection of models that were not originally presented as such).

There can however be some misunderstanding concerning the case of completely random group formation, or, in other words, the absence of assortativeness. On page 138 of BCH, one can read: "When the variance among groups is zero, A's no longer have the advantage of being in groups with disproportionally many A's. In this case group selection is inoperative, so only a costless form of group benefit could proliferate". The point where $\mathbb{P}\left(A \mid A\right) - \mathbb{P}\left(A \mid N\right) = 0$ therefore seems to be associated with the point where the (sample) variance among groups is 0.[2]

---

[2] See Van Veelen (2005) for the use of terms like variance, sample variance, expectation and

This is not correct; also with random group formation, group composition will vary from group to group. If we for simplicity think of groups of size 2, $p = 0.5$ and an infinite (or very large) population, then random group formation leads to (approximately) 25% AA groups, 50% AN groups and 25% NN groups. So here the variance among groups is not 0, but the probability of being paired to an A-player is 0.5 both for A's and for N's. It is therefore important to realise that, as can be deduced from Hamilton (1975), it is *excess* variability that is needed to select (strong) altruism, that is, more variability than the variability that comes with random group formation. In case of random group formation there is also variability, but A-players on average face the same composition of the rest of the group they are in as N-players do. However, if the variability increases from there, then the probability for A-players meeting A-players and the probability for N-players meeting A-players start diverging and if the size of this gap is big enough, given costs and benefits, A-players have a selective advantage. The point where variance among groups is 0 - which with group size 2 and $p = 0.5$ would be a population that consists of 100% AN groups - would in fact imply anti-assortative matching, and in this particular example it would come with $\mathbb{P}(A \mid A) - \mathbb{P}(A \mid N) = -1$. Contrary to what the quote suggests, even a costless form of group benefit could not proliferate here, but spiteful behaviour could, if the cost inflicted on the other outweighs the cost bore by oneself.

In BCH the working of the standard group selection model is illustrated with the following figure.
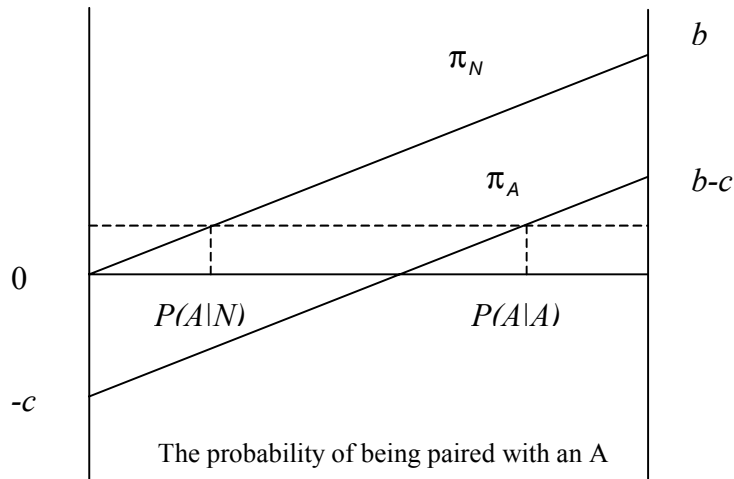

Fig. 1 from BCH, page 138.

In this figure, π_A is the payoff of an A-player as a function of the probability of being paired with an A-player, and π_N is the payoff of an N-player, also as a function of the probability of being paired with an A-player. The message of the picture is that for (strong) altruism to be selected, the difference between those probabilities, conditional on the type that the player itself is, must be large enough.

3.1 The first stages of the simulation model

As the standard group selection model does not involve groups engaging in conflicts, it can be useful to first figure out how the simulation model without war translates to this setting. What can be slightly confusing, is that fitness cost and benefits in the standard group selection model in BCH are given by c and b, while the same letters are used for parameters in their simulations. This might induce readers to think that 1) performing the act is (strongly) altruistic by definition, as $c > 0$ and $b > 0$, and that 2) by definition it implies an efficiency gain within the group, since $b > c$. It follows from the translation of the first two phases of the simulations that the latter is not true. The first is shown not necessarily to be true in section 4.

Because group size is fixed by assumption, the actual game within the group is a zero-sum game; the number of offspring of the group as a whole does not change, so all a person can change is the distribution of the odds within the group. The payoffs determine the shares of individuals in the probability distribution from which the offspring is drawn. The next generation's group members are drawn one at a time, and every time the probability that it is your offspring equals your payoff divided by total payoffs in the group. What seems to be an efficiency gain ($b > c$) therefore really is not, because probabilities have to add up to one and are normalized by dividing individual payoffs by total payoffs. The behaviour therefore not only decreases the fitness to the actor and increases the fitness of the actual recipient, but it also decreases the fitness of all group members - including donor and recipient - by increasing the denominator. Given that we know that total fitness benefits and total fitness costs are equal, A-players only shift a certain fixed amount of fitness away from themselves and divide it, in expectation, equally among the other group members, without any efficiency increase. In that case, Hamilton's rule predicts that only if groups form completely assortative, so that there are no mixed groups and $\mathbb{P}(A \mid A) - \mathbb{P}(A \mid N) = 1 - 0 = 1$, altruism will not be selected against. It will be selected against as soon as it is smaller than 1 and the picture for the simulation model, excluding the war phase, is as follows:
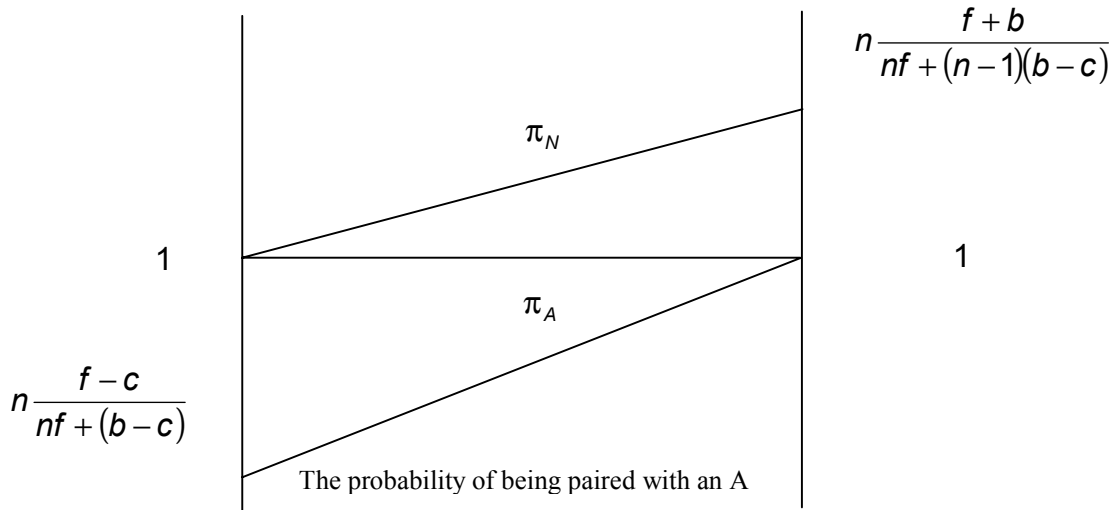
Fig. 2. Payoffs in the group phase of the simulations.

Again, π_A is the payoff of an A-player as a function of the probability of being paired with an A-player, π_N is the payoff of an N-player also as a function of the probability of being paired with an A-player, f is baseline fitness (which is 10 in BCH) and n is the group size. The formulas for calculating fitness in the picture are explained in Section 4, but the most important thing to realise is that an altruist surrounded by altruists and a non-altruist surrounded by non-altruists both have the same fitness, which is 1 here. Equality can therefore only be achieved if $P(A \mid A) = 1$ and $P(A \mid N) = 0$. The line π_A is a little steeper than π_N because if there are many altruists already, the altruistic act confers less fitness to the others, since the total, unnormalized group payoff is larger. The difference between A-players and N-players therefore is smaller at $P = 1$ - the right side of the figure - than it is at $P = 0$.

An obvious conclusion is that without the war phase, A-players could never be selected for. This conclusion is still valid if we allow for norms to evolve. To see why, we first have to see that both norms work on this picture in the same way; both a tax t and a segregation parameter s reduce the slope of the two lines. The fitness of an N-player with $\mathbb{P}(A \mid N) = 0$ however will remain at 1 and the same holds for an A-player with $\mathbb{P}(A \mid A) = 1$. This indicates that however strong a norm is, without war this simulation model could not result in selection of altruism, whatever the variability in group composition. More precisely: only for $s = 1$, $t = 1$ or $\mathbb{P}(A \mid A) - \mathbb{P}(A \mid N) = 1$ will there not be selection against A-players.

$$n\,\frac{f+b}{nf+(n-1)(b-c)}$$

$$\pi_N$$

$$\pi_A$$

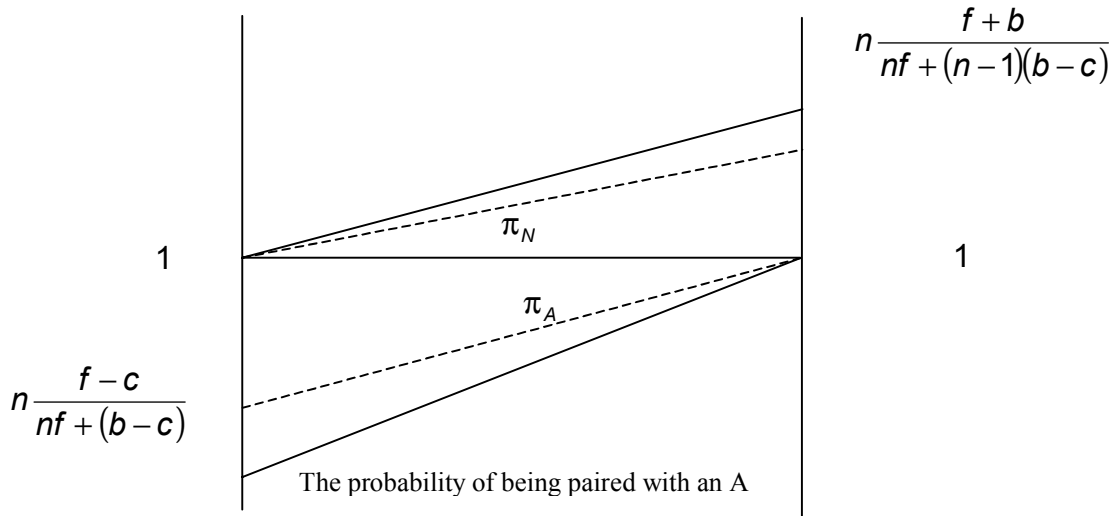$$n\,\frac{f-c}{nf+(b-c)}$$

The probability of being paired with an A

Fig. 3. Payoffs in the group phase with norms.

The dotted lines are the payoffs in the group fase of the simulations with norms. Otherwise, everything is the same as in figure 2.

In order to see how large 'relatedness' really was, we reran the simulation model. To compute it empirically, it is important to realise that it might be frequency dependent. If it is, then that could make a difference for the dynamics and it is a little problematic to speak of relatedness or of *the* difference in conditional probabilities. If we however insist on summarizing this difference and reducing it to one single number, then a relatively reasonable expression to compute is the average difference in probabilities, and that can be done by first conditioning on overall frequency $\overline{p}$ and then integrate over the distribution of frequencies as they occur in simulation runs (see Appendix A). The average difference is then as follows:

$$\int_0^1 \left(\mathbb{P}\left(A \mid A, \overline{p}\right) - \mathbb{P}\left(A \mid N, \overline{p}\right) \right) dP\left(\overline{p}\right)$$

where $\mathbb{P}\left(A \mid A, \overline{p}\right)$ is the probability of being paired with an A-player if you are an A-player yourself and the overall frequency of A-players is $\overline{p}$, $\mathbb{P}\left(A \mid N, \overline{p}\right)$ is the probability of being paired with an A-player if you are an N-player and the overall frequency of A-players is $\overline{p}$, and P is the probability measure that represents the distribution of frequencies $\overline{p}$ during simulation runs. We found that it is 0.085 for default values (same group size).

4 Reproductive externalities

In Weibull & Salomonsson (2006) reproductive externalities are introduced in a situation with random matching and finite symmetric games. This is not exactly the setting of BCH, where groups stay together and are only stirred by migration, rather than break up and reform randomly every period. Yet it can help clarify what reproductive externalities are and what they can do, if we look at how this game would be analysed in a random matching setting. More importantly, it also shows that with the game from the simulations, the (seemingly) altruistic behaviour can be selected for without any excess variability, which drives selection in the standard group selection model. In passing it also gives a nice and interesting example of a frequency dependent version of Weibull and Salomonsson's model. So while the standard group selection crucially depends on assortative matching, we now take exactly that aspect out of the simulation model by assuming that every generation, groups are completely reformed at random. Why the remaining features still make it a candidate for the label "group selection model" (which is how Weibull and Salomonsson qualify it) will be discussed in Section 5.

Suppose that groups all have the same size n, the population is infinite, and all players can be A or N. If an A-player faces i other A-players in her group, then her expected payoff is $(f + b\frac{i}{n-1} - c)/(f + (b-c)\frac{i+1}{n})$. The numerator is baseline fitness f, plus the benefit b times the probability with which it will be matched with an A-player, minus the costs c she bears because she is an A player. The denominator is the average payoff in the group and it thereby normalizes the expected payoffs because group size is fixed. If an N-player faces i A-players in her group, then her expected payoff is $(f + b\frac{i}{n-1})/(f + (b-c)\frac{i}{n})$.

In Weibull and Salomonsson these payoffs translate directly to a derived game, but in our setting this derived game is frequency dependent. The derived payoff of the average A-player, that is, its expected number of offspring after a possible war, depends on the frequency p and can be computed as follows. The probability of being in a group where i of the others are A-players is $\binom{n-1}{i}p^i(1-p)^{n-1-i}$. The probability of being matched, as a group, to a group with j A-players is $\binom{n}{j}p^j(1-p)^{n-j}$. Groups with more A-players than the other group get to replace the other group, and hence double their payoffs, groups with less see their payoffs reduce to 0, and groups that draw just keep their payoffs. With a probability k of war, the expected payoff of an A-player is

$$\sum_{i=0}^{n-1}\binom{n-1}{i}p^i(1-p)^{n-1-i}$$

$$\frac{f+b\frac{i}{n-1}-c}{f+(b-c)\frac{i+1}{n}}\left\{(1-k)+k\left(2\sum_{j=0}^{i}\binom{n}{j}p^j(1-p)^{n-j}+\binom{n}{i+1}p^{i+1}(1-p)^{n-(i+1)}\right)\right\}$$

10

while the expected payoff of an N-player is

$$\sum_{i=0}^{n-1} \binom{n-1}{i} p^i \left(1-p\right)^{n-1-i}$$

$$\frac{f+b\frac{i}{n-1}}{f+(b-c)\frac{i}{n}} \left\{(1-k) + k\left(2\sum_{j=0}^{i-1}\binom{n}{j}p^j\left(1-p\right)^{n-j} + \binom{n}{i}p^i\left(1-p\right)^{n-i}\right)\right\}$$

With the benchmark values of BHC, that is $n = 20, f = 10, b = 2, c = 1, k = 0.25$, the expected payoffs of the two strategies as functions of the frequency p are plotted in the following picture.
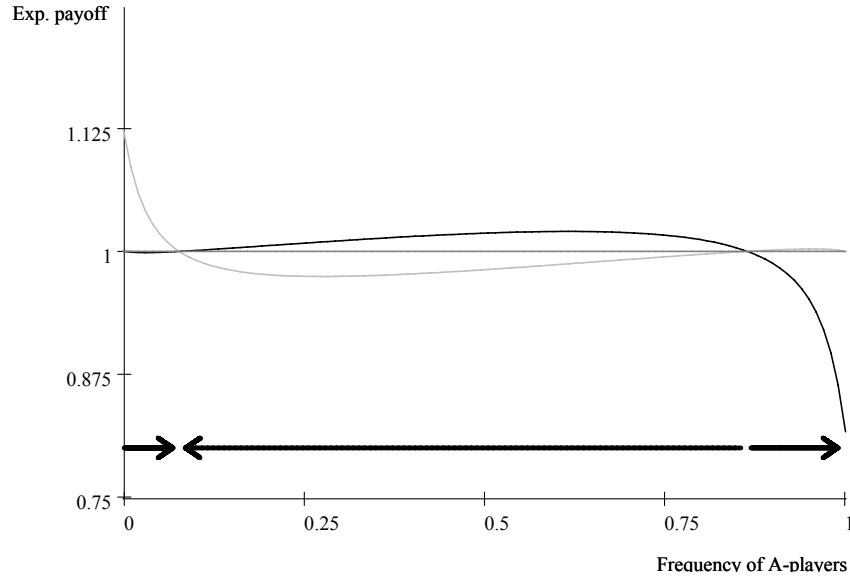

Fig. 4. Expected payoffs A's (grey line) and N's (black line) for benchmark values.

The game has, from left to right, a stable mixed equilibrium, an unstable mixed equilibrium and a stable pure equilibrium. For different values of n and k this picture changes accordingly, with payoffs for A-players going down if the group size goes up, and going up along with the probability of wars.

As an illustration of the dynamics, it can be useful to consider what happens at both extremes. At $p = 1$, that is, in a population that consists of A-players only, any individual that mutates into an N-player makes the group it is in lose a war with probability 1, against a probability of winning of 0.5 had it not mutated. So if c is not too large, selection at $p = 1$ favours A-players, which makes it a stable equilibrium. At $p = 0$, that is, in a population that consists of N-players only, any individual that mutates into an A-player makes the group it is in win a war with

11

probability 1, against a probability of winning of 0.5 had it not mutated. So if c is not too large, selection at $p = 0$ favours A-players, from which we can conclude that it can not be an equilibrium.

It can be worthwhile to realise that at both equilibria in the above picture the behaviour simply optimizes the fitness of the individual and no arguments concerning who receives the benefits of the so-called altruistic act play a role. The reproductive externalities here consist of the increased probability of surviving a possible war. While it saves costs c to be an egoist, it lowers the odds of surviving a conflict situation with other groups. In the classification of Lehmann & Keller (2006), this model would therefore fall within the category of direct benefits. The outcome also shows that excess variability, unlike in the standard group selection model, is not needed to get selection of A-players; both stable equilibria have A-players, even though groups are formed randomly. Whether or not we qualify the A-players as altruistic is perhaps a matter of taste. If we define (strong) altruism as behaviour that lowers an individual's fitness, relative to the whole population, then it depends on the current frequency whether or not it is (strongly) altruistic. In equilibrium however it is not. This is an unambiguous, clear and possibly preferable definition of altruism. Another option is to restrict the horizon or the scope of the definition to times of peace. In that case we would simply always include this A-player's behaviour, that may very well match with what in everyday life passes the test of genuine altruism. Furthermore, even in equilibrium, the behaviour of A-players does always satisfy Wilson's (1979, 1990) definition of weak altruism; in equilibrium, A-players do not decrease their fitness compared to the whole population by playing A, but they do increase their group members' fitness, so the other group members gain more from the behaviour of an A-player than it does itself.

Given that we can find out whether or not A-players have a selective advantage by looking at direct benefits only (see Lehmann & Keller, 2006), one may wonder why this is nonetheless called a group selection model. This will be discussed in Section 5, but for now it is important to see that we have only taken the game from the simulations, eliminated the process that makes group formation assortative, and find that A-players can still be selected for.

It is also worth noting that of the two composing elements of the simulation model, the second one - tying together the fate of group members - is the innovative ingredient. Although BCH do follow Darwin (see Section 5), conflicts between groups were, to our knowledge, not modelled explicitly before in this context.

4.1 Evolution of norms

The norms that feature in BCH can evolve because of a combination of ingredients from both models. First we should note that within the standard group selection

model, norms could also evolve together with altruism. We then would need to assume that groups can indeed grow in size, which is not the case in the simulations, but which is essential for the evolution of altruism as well as norms within this model (see Section 3). Furthermore we would have to assume that groups are not reshuffled every stage in an assortative way, but that there is a group fissioning process that causes relatedness within groups, and that causes group norms also to be conserved to some extent during a group's history. This does happen in the simulations. Groups with higher t or s then would grow faster, or larger, than groups with lower t or s. The reason is that higher norms have a positive effect on the share of A-players in the group, and groups with more A-players grow faster. So if not counterbalanced by some cost of the norm, higher values of t and s would always spread at the expense of lower values.

In the second type of group selection model, where groups are randomly reshuffled every period, it is a little hard to imagine how norms are retained within groups, as groups only exist for one generation. If that could be overcome, and higher levels of norms can cause higher frequencies of A-players within groups, then there would be a very good reason for norms to spread, as groups with more A-players win wars with higher probability.

The actual simulations combine the ingredients from both models in such a way that norms can evolve. Although groups cannot grow in size - thereby violating a necessary condition in the pure standard group selection setting - the groups are only marginally changed between rounds and therefore correlation between the fraction of A-players and the level of the norms can build up. Combined with an ingredient from the other type of model - where groups with many A-players win wars more frequently - this allows the norms to spread.[3]

---

[3]The simulations have two norms, a sharing norm t and a segmentation norm s. Although the sharing norm is perhaps more appealing, it is also slightly problematic for the following reason. In the model under consideration, the only individual trait on which selection act is altruism. Given this model, it is clear that the presence of a norm reduces the selection pressure within a group against altruists and thereby increases the odds of surviving inter-group conflicts. This norm however not only reduces selection pressure against altruists, it also reduces selection pressure against a lot of other things. Lazy farmers for instance and incapable hunters benefit from it as well. And so genes for altruism thrive along with genes for laziness and low hunting ability, whereby the first may increase the total payoff of the group, while the other two decrease it.

A possible solution for this problem may lie in interaction with sexual selection. If benefits from efficient farming and good hunting also lie in the increased number of interested mates, then a sharing norm reduces the selection pressure against lazy farming and bad hunting less than the pressure against altruism. Once a preference for good farmers or hunters is established, the survival value of it is not needed to sustain the feedback loop of trait and preference for the trait.

4.2 Combining the two

The simulation model of BCH combines the ingredients of both types of models. Group formation is not random and during the simulations relatedness builds up. (How much depends on the parameters, but there is always some.) The war phase on the other hand introduces reproductive externalities and causes the interests of group members to align to some extent. (Again, to what extent depends on the parameters). Together, these two ingredients make a simulation model in which altruism as well as norms spread for some and do not spread for other parameter values. Even if relatedness is not extremely high (it is for instance 0.085 for default values) then still the benefit of increasing the probability of surviving a war is shared by the whole group, members of which are slightly related to the actor. For the default group size of 20, this amounts to an additional contribution to the expected number of A-players surviving the group conflict phase that can be roughly approximated by $20 \times 0.085 = 1.7$ times the value of the increase for the individual itself. Note that, first of all, benefits are not additive, and also that benefits and costs can not be added to / substracted from each other. If benefits $\widetilde{b}$ are measured as increases in survival probability and costs $\widetilde{c}$ as reductions in expected number of offspring conditional on survival, then fitness becomes $\left(1 + \widetilde{b}\right)\left(1 - \widetilde{c}\right)$.

It is furthermore important to realise that benefits and costs are no longer exogenously given values, but endogenous variables, values of which not only are frequency dependent but also depend on the choice of parameter values. So even though the behaviour under study itself is constant across different simulations, it is the choice of parameters and the overall frequency that determine how altruistic the action is.

5 Telling the different types apart

It can be useful to summarize the differences between the two types of models and relate that to the way they are treated in the literature. In a book that strongly advocates group selection theory (Sober & Wilson, 1998) the distinction between the types is not made - or at least not explicitly - but both of them are present. The first (or standard) type is discussed most elaborately (pp 17-100) and seems to define their idea of group selection. As we have seen, the defining feature there is that altruists are more likely to interact with (other) altruists than egoists are. Whichever way the accounting is done, it follows that assortativeness in group formation allows for altruism to evolve. From the individual point of view the argument is as follows: altruists incur more costs than egoists because they are altruists. They also enjoy more benefits than egoists do, because they interact more with altruists. Whether a given form of altruism can evolve, depends on how

14

these costs and benefits are balanced. From the 'gene for altruism' point of view the end result is the same; an altruistic act is of course costly, but the benefits are enjoyed by a disproportional share of altruists, which then may do more good than harm for the share of altruistic genes in the total population.

Their explanation of how increased variability in group composition can lead to altruism is accompanied, however, with a discussion of frameshifting the idea of hierarchical selection from *'gene-individual'* to *'individual-group'*. Group selection is then to be seen as another major transition in evolution, like the transition from single celled to multicellular organisms. Here it is important to realise that the explanation of why cooperation is stable for different cells within one body is *not* that they share 100 % of their genes. Even though it is very well possible that full relatedness has played a role in the transition itself, the reason why those cells keep on cooperating lies in the fact that they are sharing one and the same vehicle of reproduction. As a result of that, their interests are, on the relevant time scale, very much in line. The dilemma that a cell faces is therefore not whether or not it will increase the offspring of the other cell at the cost of reducing its own, but the question is to produce more offspring together or not. If one wants to force this into a game theory setting, then the appropriate game is not a social dilemma, but a mutualistic game, as it is not against the cells own reproductive interest to cooperate.

The parable of the boat race in Dawkins (1976, p38-39) explains how genes that cooperate in making one working organism can be selected for. Sober & Wilson (1998, p97, p124) apply the same reasoning to individuals and groups. They are aware of the fact that there may be limits to the analogy, but their assessment of to what extent interests of individuals within a group can align is nonetheless relatively positive. Although one can have reservations concerning the extent of the alignment of interests, it is undeniably true that in the simulation model of BCH fates of individuals at least coincide a bit. In times of peace they may compete for resources, but a war is survived together or not at all. It is therefore important to identify this aspect of the simulation model as distinct from the assortative matching idea and it is worthwhile realising that the two different types of group selection models have different reasons why they work, if they do. The working of the first one depends on the answer to the question *'is the probability that you also have the gene large enough'*[4] and the other needs an affirmative answer to *'are our interests enough in line'*. In many situations the answer to both questions may be the same, but it is important to realise that either model can work without the ingredient that is crucial for the other. Also in the actual simulations there is a

---

[4]If we allow for assortment of phenotypes that not necesarily implies assortment of genotypes one could also write this as *'is the probability that you are an altruist too large enough'*.

bit of both, but we hope we have convincingly argued that for this particular case the behaviour can not be selected for in the pure standard group selection model (that is, without warfare and with groups of constant size), while *A*-players can be selected in the pure reproductive externalities case (that is, without assortative matching). In the simulations the warfare part therefore is every bit as important as the relatedness that the process induces.

5.1 Different labels & terminology

In the literature, both models are referred to in different ways. The first type of model is called a group selection model by a variety authors, including Sober & Wilson (1998) and Bowles, Choi & Hopfensitz (2003). Other authors, such as Lehmann & Keller (2006) and Grafen (2007) classify this type of model as a kin selection model. The second type of model is called a group selection model by, again, Sober & Wilson (1998), while others, such as Lehmann & Keller (2006), would classify it as a model with direct benefits. We do not have a preference for a specific set of labels for these two types of models. What we do think is important is that they are clearly distinguished, and hence that those who prefer to refer to both types as group selection models, can improve clarity by adding an indication of the type. It is also important to realize that the insight that group selection and kin selection are different but equivalent ways to describe the same process (see for instance Lehmann et al., 2007 and Grafen, 2007) applies to the first type of group selection model, and not to the second.

5.2 Racial discrimination or groupishness?

Humans behave nicer towards group members than they do towards individuals that do not belong to their group. In *Innate social aptitudes of man* Hamilton (1975) chooses what we call the standard group selection model in order to explain inter-group hostility and intra-group friendliness. Relatedness of an individual to out-group members is lower than to in-group members, hence the difference in behaviour. Experimental evidence, starting from Tajfel & Turner (1986) has shown that the difference in behaviour is still there if individuals are randomly assigned to groups. In that case the actual difference in relatedness and therefore the supposed evolutionary reason is no longer there, but the behaviour remains. One explanation within the standard group selection model could be that in our evolutionary past humans hardly ever were in a situation where group members were not related (that is, group members in expectation shared more genes with each other than with the average population member). In that case humans are just not adapted to a situation with randomly composed groups.

In the other group selection model, however, being in the same group is not just a cue for relatedness. If an individual is in the same group, then that is something that matters by itself; it implies that the fate of surviving a possible conflict is shared. Here relatedness is not an issue, so it is not maladaptive to behave nicer to group members than to non-group members even if the groups are formed randomly.

Going for the first type of model, Hamilton (1975, p134) writes

> I hope to produce evidence that some things which are often treated as purely cultural in humans—say racial discrimination—have deep roots in our animal past and thus are quite likely to rest on direct genetic foundations. To be more specific, it is suggested that the ease and accuracy with which an idea like xenophobia strikes the next replica of itself on the template of human memory may depend on the preparation made for it there by selection - selection acting, ultimately, at the level of replicating molecules.

His explanation implies that (standard) group selection has been at work and racial discrimination is expected to be as strong as relatedness is.

In *The origins of virtue*, Ridley (1997, p188) on the other hand claims that it is not the interests of others that share the genes for altruism, but a keen eye for our own individual interests that shape group-behaviour. His conclusion is that this refutes group selection as a mechanism.

> The fact that people form emotional attachments to groups, even arbitrary ones, such as randomly selected sports teams, does not prove group selection, but the reverse. It proves that people have a very sensitive awareness of where their individual interests lie - with which group. We are an extremely groupish species, not a group-selected one. We are designed not to sacrifice ourselves for the group but to exploit the group for ourselves.

Here Ridley seems to claim that finding individuals being groupish would exclude group selection. It is interesting to note that in the Weibull and Salomonsson case with reproductive externalities, individuals do indeed behave according to their individual interests. Nonetheless they classify it as a group selection model. It is a model that is rather different from the hierarchical selection models that we are used to and that depend on relatedness or assortative matching, but they are not alone in their choice. Many authors quote the following excerpt from *The descent*

of man, Darwin (1871, p166) to argue that he thought group selection a possibility, and the process he describes here does coincide with the second type of model:

> It must not be forgotten that although a high standard of morality gives but a slight or no advantage to each individual man and his children over the other men of the same tribe, yet that an increase in the number of well-endowed men and an advancement in the standard of morality will certainly give an immense advantage to one tribe over another. A tribe including many members who, from possessing in a high degree the spirit of patriotism, fidelity, obedience, courage, and sympathy, were always ready to aid one another, and to sacrifice themselves for the common good, would be victorious over most other tribes; and this would be natural selection. At all times throughout the world tribes have supplanted other tribes; and as morality is one important element in their success, the standard of morality and the number of well-endowed men will thus everywhere tend to rise and increase.

In this passage, it seems evident that tribes are thought of as units that compete directly with each other, as they can 'supplant' and 'be victorious over' other tribes. Herewith it differs from the standard hierarchical selection models that for instance Price (1972), Hamilton (1975) and Frank (1998) have in mind. In those models, the payoffs determine the frequencies in and the size of subsequent groups, but there is no explicitly modelled reason why overall populations would not dwindle or explode. This can be overcome by assuming that when some constraint is binding, all payoffs should be normalized. A natural way to do this is to divide them by the sum of all payoffs of all groups and multiply with the overall carrying capacity. Such a normalization keeps the total population constant, but its form rests on the assumption that every member of the total population competes with every other member equally intensely for a place in the next generation. In many situations that is a very reasonable assumption, especially if the life cycle includes a period spent within the group as well as a period in which groups are dissolved and all individuals live alone, competing for some scarce resource. In the simulations in BCH however there is no such life cycle. Moreover, the assumption that groups regularly interact violently gives competition between individuals a particular structure. It is perhaps an uncomfortable idea to model violent competition between groups so explicitly, but that does not make it unrealistic and it works out Darwin's idea of the evolution of morality quite literally. The result of this assumed structure of violent interactions is that there are externalities that, ironically, drive the evolution

of pro-social behaviour within the group. BCH's model therefore is a combination of a standard hierarchical selection model with a group selection model in the spirit of Darwin, and it has individuals that are groupish - more than racist - as well as group-selected.

5.3 Searching for the right relatedness

Hamilton's rule in principle provides a testable implication of the standard group selection model. If fitness costs and benefits are known, we have to measure relatedness (differences in conditional probabilities) to see if it is high enough to explain the altruistic behaviour. The model with reproductive externalities on the other hand does not need relatedness to work. If that were the true model, and we misinterpret c and b in this model as actual fitness costs and benefits, then we find that relatedness is 0 and therefore reject the standard group selection model. What goes wrong is that we have overlooked the reproductive externalities that makes actual fitness costs 0 and not c.

More important is that the simulations suggest that the true model might very well be a mixture of the two. If we again (mis)take c and b for actual fitness costs and benefits, we will find that relatedness is positive, but still too low to explain altruistic behaviour with the standard group selection model only. It is important to note that the measurement issue of the true fitness costs and benefits is at least as big a problem as the measurement of relatedness. As we have seen in subsection 4.2, fitness costs and benefits of a given behaviour have become endogenous variables in the model, and will differ between different choices of parameters in the simulation. As the behaviour itself is the same across all those choices of the parameters in the simulation model, and benefits may not even be reaped within one generation, it is not strange that measuring the actual fitness costs and benefits is rather hard. An interesting and recent empirical study in this respect is Bowles (2006).

6 Markov chains

6.1 A simple example

Apart from distinguishing the different composing elements, it can also be useful to recognise the basic structure of the actual simulations, which is that the simulation model is a Markov chain. This is useful if we interpret simulation outcomes, and especially if we want to use the data that the simulations produce to draw conclusions concerning the dynamics. Simulation results therefore are best seen as characteristics of the dynamics in a Markov chain that will differ for different choices of parameter values. In order to illustrate what kind of Markov chain the

model is, we take the simplest possible example with only two groups consisting of only two players that can be either A or N. The six states are numbered as follows:

$$
\begin{array}{ll}
1 & (NN, NN) \\
2 & (NN, NA) \\
3 & (NN, AA) \\
4 & (NA, NA) \\
5 & (NA, AA) \\
6 & (AA, AA)
\end{array}
$$

One whole round (or generation) can be decomposed in matrices for every stage; K for the war phase, R for the reproduction, E for the mutation and M for the migration. The matrices are given in Appendix B. A whole round, starting with war and ending with migration, is represented by the product matrix $MERK$. The invariant distribution p^* is the solution of $MERKp = p$ and for benchmark values $k = 0.25, f = 10, c = 1, b = 2, \epsilon = 0.001$ and $m = 0.1$ it is, numerically:

$$
p^* = \begin{bmatrix}
0.088 \\
0.001 \\
0.005 \\
0.001 \\
0.006 \\
0.899
\end{bmatrix}
$$

The average frequency of altruists in the invariant distribution is the inner product of the vector v that consists of frequencies of altruists in the different states - so here $v = [0, 0.25, 0.25, 0.5, 0.75, 1]^T$ - and the invariant distribution p^*, which makes $\langle v, p^* \rangle = \sum_{i=1}^{6} v_i p_i^* = 0.9$. We can also compute the expected change in frequency of altruists for every state by taking the difference $[MERK]^T v - v$. The i'th component of this vector is the expected change in frequency of A-players when in state i. In this case, that is:

$$
[MERK]^T v - v = \begin{bmatrix}
0.001 \\
0.018 \\
0.125 \\
-0.071 \\
0.035 \\
-0.001
\end{bmatrix}
$$

The combined effect of the four matrices can be described as follows. Mutations create a very small outflow from (NN, NN) and (AA, AA). From (NA, NA) the expected change in the overall fraction of altruists is negative, which is to be expected, because wars can not increase it, but within groups A-players in expectation do worse than N-players. The other states show an expected increase in the overall fraction of A-players

To compare this to the case with random matching in an infinite population, we below plot the payoffs of altruists and non-altruists with $N = 2, b = 2, c = 1, k = 0.25$. The only equilibrium is pure; all play A. Interestingly, the simulations will, by contrast, show on average a proportion of 8.8% of non-altruists. This is caused by the relative stickiness of both (NN, NN) and (AA, AA), from which the only way out is by mutation. Although the tendency for the share of A-players is up, chance might bring the Markov chain to (NN, NN), from which it takes on average a very long waiting time to leave.
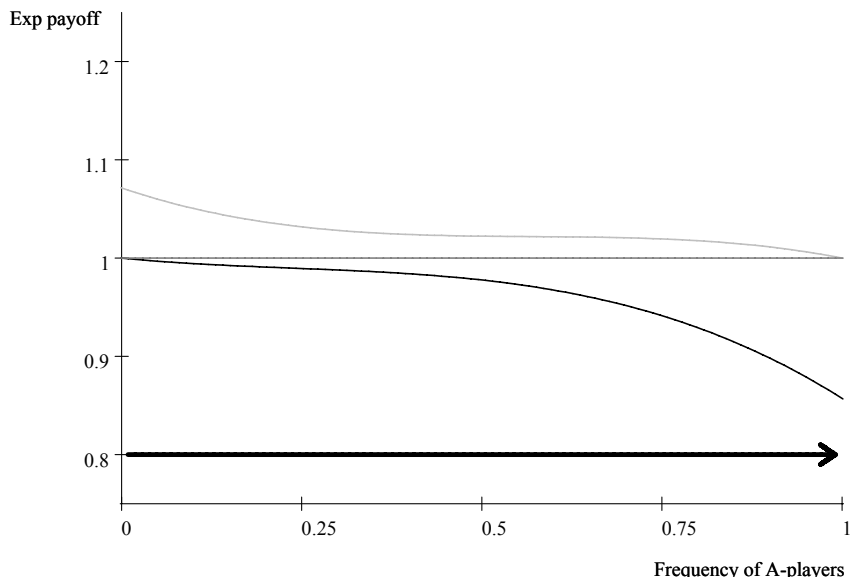

Fig. 5. Expected payoffs A's (grey line) and N's (black line)

6.2 Statistical estimation of the Price equation

The actual simulations are also Markov chains, but there is a complication that concerns the size of the state space. If we take default values of 20 groups and a total of 400 individuals, then even if we restrict ourselves to situations where all groups have equal size, the number of states[5] equals $\binom{40}{20} = 137\,846\,528\,820$. The

---

[5] A state is then characterized by 20 numbers, that represent how many groups with 20 A-players there are, how many with 19 A-players, and so forth. Distributing 20 groups over 20 positions can be done in $\binom{20+20}{20}$ ways.

state space is therefore best described as enormous and without further information on transition probabilities, simulation runs of feasible length would be insufficient to make any statement about the average frequency of altruists in the invariant distribution. In this simulation model, however, the different transition probabilities are related to each other, as they are the result of the same type of probability calculation with only differing number inputs that correspond with the two relevant states. Thus there is reason to expect that even when two consecutive runs of the simulation hardly share states they both have visited, the dynamics may be very similar in terms of frequencies of A-players. It is therefore natural to look for an overall characterization of the probabilities of changes in the overall frequency of altruists as a function of simple characteristics of the states, such as current frequency of altruists and some measure of the variability of the frequency of altruists within the different groups. The regressions in section 3 of BCH can be viewed as a way to try to do this. In the paper they are described as a statistical estimations of the Price equation, and although we now know from Van Veelen (2005, Section 6) that applying the Price equation as a statistical device to group selection models is generally not such a good idea, we can still ignore the name or what inspired it and simply look at the regression itself to see whether it could be informative about the dynamics.

The dependent variable in the regression is the change in the overall frequency of A-players - $\Delta \overline{p}$ - and the explanatory variables are a measure of within group variability - $\sum_i p_i (1 - p_i) \frac{n_i}{N}$ - and between group variability - $\sum_i (p_i - \overline{p})^2 \frac{n_i}{N}$ - where p_i is the frequency of A-players in group i, n_i is the number of individuals in group i, $N = \sum_i n_i$ is the total number of individuals and $\overline{p} = \sum_i p_i \frac{n_i}{N}$ is the overall frequency of A-players. Indexing all observations with $l = 1, ..., L$, this suggests that dependent and explanatory variables are related in the following way

$$\Delta \overline{p}_l = \alpha + \beta_1 \sum_i p_{i,l} (1 - p_{i,l}) \frac{n_{i,l}}{N} + \beta_2 \sum_i (p_{i,l} - \overline{p}_l)^2 \frac{n_{i,l}}{N} + \varepsilon_l$$

where ε is some iid variable with mean zero. If we restrict attention to the mean of the error terms only, this would reduce to the assumption that

$$\mathbb{E} \Delta \overline{p}_l = \alpha + \beta_1 \sum_i p_{i,l} (1 - p_{i,l}) \frac{n_{i,l}}{N} + \beta_2 \sum_i (p_{i,l} - \overline{p}_l)^2 \frac{n_{i,l}}{N} \quad \forall l$$

Because we actually know what the true model is, it can be checked whether or not this assumption is correct and it turns out that it is not. Appendix C contains a simple choice of a state space combined with BCH's benchmark values that shows that the expected change in overall frequency of altruists is not a linear function of the two proposed variables already for a very simple case. This implies that we fit a misspecified model to the data, which results in residuals that contain not only

errors induced by the transition probabilities being random variables, but also the differences between $\mathbb{E}\Delta\overline{p}$ and the right hand side of the above equation.

There is also another angle from which we can directly see that this approach cannot work. Therefore it is important to first note that BCH make *one* such estimate of the Price equation, using data generated by four simulation runs with presumably the same parameter values k, ϵ, m and N. The outcome therefore only applies to this one choice of parameter values; with different parameter values we would get different estimates. The reason is that expected changes in $\overline{p}$ in any particular state will differ across different choices for the parameters. For the incidence of war for instance, the general pattern for any single state is that it will have an expected decrease in $\overline{p}$ for low values of k and an expected increase for high values of k. And while the different k's lead to different $\mathbb{E}\Delta\overline{p}$'s, the supposedly explanatory variables stay the same, as within and between group variability both are determined by state characteristics only and not by the parameters of the simulation. This implies that such regressions have very limited meaning; $\mathbb{E}\Delta\overline{p}_l$ is not adequately explained by within- and between group variability.

The "statistical estimation of the Price equation" therefore fits a misspecified model to the data and interpretation of the regression result can lead to erroneous claims. The structure within the large state space of the Markov chain however still remains, and therefore we think that it is justified to rely on simulation runs of otherwise insufficient length to represent characteristics of the invariant distribution. It therefore seems justified to compute relatedness as we did above, or use output data to compute any averaged characteristic we are interested in as a function of the parameters. The dependence of such characteristics on the parameters however is essential, as most of the things we are interested in, such as for instance fitness differentials, are endogenous

7   Final remarks

We hope we have convincingly argued that the simulations in BCH combine ingredients from two different types of model that on their own have quite different implications. The first could be called the standard group selection model and here assortative group formation is essential for selection of altruism. In the second type of group selection model the interests of group members are to some extent aligned. Understanding and distinguishing the exact workings of these different ingredients of the simulation model we think is useful, because both of them are likely to have played a role in the evolution of human altruism and norms. While the first type has received quite some attention in the literature, compared to the second one, we think that the second is no less relevant.

Making this distinction is a contribution to the effort made recently by evolutionary biologists to clarify the significance of published papers (see Grafen, 2006, Lehmann et al, 2006) and classify them according to what drives the results (Lehmann & Keller, 2006). This is an important thing to do in a field that is facing a plethora of models, in which the mechanism at work is not always explicitly identified. In the classification of Lehmann & Keller, the first type of group selection model can be understood as a kin selection model, while the second type can be understood in terms of direct benefits.

The last section of the paper, that focusses on how to handle output data in simulations that are Markov chains, follows Van Veelen (2005) in suggesting that statistics and probability theory are easily confused, and this should carefully be avoided when interpreting simulation results. These simulation models can, as a rule, be described as Markov chains and we should aim at describing the results as characteristics of invariant distributions.

8   Acknowledgements

References

[1] Bowles, S., 2006. Group competition, reproductive leveling, and the evolution of human altruism. *Science* **314**, 1569-1572.

[2] Bowles, S. Choi, J-K. & Hopfensitz, A., 2003. The co-evolution of individual behaviors and social institutions, *Journal of Theoretical Biology* **223**, 135-147.

[3] Darwin, C., 1871. *The descent of man, and selection in relation to sex.* London: John Murray

[4] Dawkins, R., 1976. *The selfish gene.* Oxford: Oxford University Press.

[5] Frank, S.A., 1998. *Foundations of social evolution*, Princeton University Press, Princeton, New Jersey.

[6] Grafen, A. 2007. Detecting kin selection at work using inclusive fitness. *Proceedings of the Royal Society B* **274**, 713-719.

[7] Hamilton, W.D., 1975. Innate social aptitudes of man: Approach from evolutionary genetics. In Biosocial Anthropology (Fox, R., ed.) 133-155. New York, Wiley.

[8] Lehmann, L & Keller, L., 2006, The evolution of cooperation and altruism - a general framework and a classification of models. *Journal of Evolutionary Biology* **19**, 1365-1376

[9] Lehmann L, L. Keller, S. West & D. Roze., 2007. Group selection and kin selection: two concepts but one process. Proceedings of the National Academy of Sciences **104 (16)** 6736-6739.

[10] Price, G.R., 1970. Selection and Covariance. *Nature* **227**, 520-521.

[11] Price, G.R., 1972. Extension of covariance selection mathematics. *Annals of human genetics* **35**, 485-489

[12] Ridley, M. 1996. *The origins of virtue: human instincts and the evolution of cooperation.* London: Viking.

[13] Sober, E., Wilson, D.S., 1998. *Unto Others; the evolution and psychology of unselfish behaviour.* Cambridge MA, Harvard University Press.

[14] Tajfel, H and J.C. Turner, 1986. The social identity theory of intergroup behaviour. In Worchel, S & W.G. Austin (Eds.) *Psychology of intergroup relations* (2nd ed., pp 7-24). Chicago: Nelson-Hall

[15] Van Veelen, C.M., 2005. On the use of the Price equation, *Journal of Theoretical Biology* **237**, 412-426.

[16] Weibull, J.W., Salomonsson, M., 2006. Natural selection and social preferences, *Journal of Theoretical Biology* **239 (1)** 79-92.

[17] Wildschut, T., B. Pinter, J.L. Vevea, C.A. Insko, J. Schopler, 2003. Beyond the group mind: A quantitative review of the interindividual-intergroup discontinuity effect. *Psychological Bulletin* **129 (5)**, 698-722.

[18] Wilson, D.S., 1979. Structured demes and trait-group variation, *The American Naturalist*, **113 (4)**, 606-610.

[19] Wilson, D.S., 1990. Weak altruism, strong group selection, *Oikos*, **59 (1)**, 135-140

A Computing relatedness

In computing assortativeness of group formation, we keep in mind that we are after conditional probabilities that have the form

$$\mathbb{P}\left(C \mid D\right) = \frac{\mathbb{P}\left(C \cap D\right)}{\mathbb{P}\left(D\right)}$$

which here means that

$$\mathbb{P}\left(\text{being matched with an } A\text{-player} \mid \text{being an } A\text{-player}\right) = \frac{2\mathbb{P}\left(AA\right)}{2\mathbb{P}\left(AA\right) + \mathbb{P}\left(AN\right)}$$

$$\mathbb{P}\left(\text{being matched with an } A\text{-player} \mid \text{being an } N\text{-player}\right) = \frac{\mathbb{P}\left(AN\right)}{\mathbb{P}\left(AN\right) + 2\mathbb{P}\left(NN\right)}$$

If i is the number of A-players in the population, we condition on the frequency $\overline{p} = \frac{i}{N}$, and compute the empirical difference in conditional probabilities for every different frequency before summing them with the appropriate weights: $(*)$

$$\sum_{i=0}^{N}\left(\frac{2 \cdot \#\left(AA, \overline{p} = \frac{i}{N}\right)}{2 \cdot \#\left(AA, \overline{p} = \frac{i}{N}\right) + \#\left(AE, \overline{p} = \frac{i}{N}\right)} - \frac{\#\left(AE, \overline{p} = \frac{i}{N}\right)}{\#\left(AE, \overline{p} = \frac{i}{N}\right) + 2 \cdot \#\left(EE, \overline{p} = \frac{i}{N}\right)}\right) \frac{\#\left(\overline{p} = \frac{i}{N}\right)}{\text{simulation length}}$$

Note that conditioning makes a difference, and that the above formula really is different from

$$(**) \qquad \frac{2 \cdot \#\left(AA\right)}{2 \cdot \#\left(AA\right) + \#\left(AE\right)} - \frac{\#\left(AE\right)}{\#\left(AE\right) + 2 \cdot \#\left(EE\right)}$$

As an example, suppose that groups are duo's, population size is infinite and group formation is random. Suppose furthermore that only two extreme population frequencies occur: $p = 0.1$ and $p = 0.9$ and that both occur with probability 0.5. Then $\mathbb{P}\left(A \mid A, 0.9\right) = \mathbb{P}\left(A \mid N, 0.9\right) = 0.9$ and $\mathbb{P}\left(A \mid A, 0.1\right) = \mathbb{P}\left(A \mid N, 0.1\right) = 0.1$. In that case

$$\int_{0}^{1} \left(\mathbb{P}\left(A \mid A, \overline{p}\right) - \mathbb{P}\left(A \mid N, \overline{p}\right)\right) dP\left(\overline{p}\right) = 0$$

and (*) will also be close to 0.

Lumping all AA-pairs together, AE-pairs together and EE-pairs together however gives a different picture; (**) will be around

$$\frac{2 \cdot \left(0.5\,(0.1)^2 + 0.5\,(0.9)^2\right)}{2 \cdot \left(0.5\,(0.1)^2 + 0.5\,(0.9)^2\right) + 2 \cdot (0.1)\,(0.9)} - \frac{2 \cdot (0.1)\,(0.9)}{2 \cdot (0.1)\,(0.9) + 2 \cdot \left(0.5\,(0.1)^2 + 0.5\,(0.9)^2\right)} = 0.64,$$

even though the probability of being matched with an A-player is equal for both types all the time. This measure therefore would incorrectly give the impression that matching is assortative.

B   Transition matrices

In order to be able to decompose the transition matrix of one round (or generation) in matrices for every stage, the timing of the different stages is altered slightly. This is inconsequential for the dynamics but very helpful for the illustration.

With probability k the two groups engage in a war. This war is won by the group with the higher number of A-players, which replaces the group with the lower number of A-players. The resulting transition matrix K is therefore

$$
\begin{matrix}
1 & 0 & 0 & 0 & 0 & 0 \\
0 & 1-k & 0 & 0 & 0 & 0 \\
0 & 0 & 1-k & 0 & 0 & 0 \\
0 & k & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 1-k & 0 \\
0 & 0 & k & 0 & k & 1
\end{matrix}
$$

where K_{ij} is the transition probability from state j to state i.

As the matching within groups is not random anymore, the payoffs are deterministic, but they do make up the probabilities with which the next generation is drawn. For baseline fitness f, costs c and benefits b, the reproduction matrix R is

$$
\begin{matrix}
1 & \left(\frac{f+b}{2f+b-c}\right)^2 & 0 & \left(\frac{f+b}{2f+b-c}\right)^4 & 0 & 0 \\[2mm]
0 & 2\left(\frac{f+b}{2f+b-c}\right)\left(\frac{f-c}{2f+b-c}\right) & 0 & 4\left(\frac{f+b}{2f+b-c}\right)^3\left(\frac{f-c}{2f+b-c}\right) & 0 & 0 \\[2mm]
0 & \left(\frac{f-c}{f+b-c}\right)^2 & 1 & 2\left(\frac{f+b}{2f+b-c}\right)^2\left(\frac{f-c}{2f+b-c}\right)^2 & \left(\frac{f+b}{2f+b-c}\right)^2 & 0 \\[2mm]
0 & 0 & 0 & 4\left(\frac{f+b}{2f+b-c}\right)^2\left(\frac{f-c}{2f+b-c}\right)^2 & 0 & 0 \\[2mm]
0 & 0 & 0 & 4\left(\frac{f+b}{2f+b-c}\right)\left(\frac{f-c}{2f+b-c}\right)^3 & 2\left(\frac{f+b}{2f+b-c}\right)\left(\frac{f-c}{2f+b-c}\right) & 0 \\[2mm]
0 & 0 & 0 & \left(\frac{f-c}{2f+b-c}\right)^4 & \left(\frac{f-c}{2f+b-c}\right)^2 & 1
\end{matrix}
$$

With mutation probability ϵ, the matrix of transition probabilities E can for small ϵ be approximated with

$$
\begin{array}{cccccc}
1-4\epsilon & \epsilon & 0 & 0 & 0 & 0 \\
4\epsilon & 1-4\epsilon & 2\epsilon & 2\epsilon & 0 & 0 \\
0 & \epsilon & 1-4\epsilon & 0 & \epsilon & 0 \\
0 & 2\epsilon & 0 & 1-4\epsilon & 2\epsilon & 0 \\
0 & 0 & 2\epsilon & 2\epsilon & 1-4\epsilon & 4\epsilon \\
0 & 0 & 0 & 0 & \epsilon & 1-4\epsilon
\end{array}
$$

This matrix only includes transitions of zero or one mutation.[6]

If migration pairs are formed with equal probability, this makes the following migration matrix M

$$
\begin{array}{cccccc}
1 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 1-m & \frac{m}{2} & 0 & 0 \\
0 & 0 & m & 1-\frac{m}{2} & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 1
\end{array}
$$

A whole round consists of the product matrix $MERK$.

C Expected changes in frequency

For this example, we will again have a population of 4, but now not with fixed and equal group sizes. This complicated matters, as the state space now expands to the following 19 states:

1	$(\cdot, EEEE)$	6	(E, EEE)	10	(A, EEE)	14	(EE, EE)
2	$(\cdot, EEEA)$	7	(E, EEA)	11	(A, EEA)	15	(EE, EA)
3	$(\cdot, EEAA)$	8	(E, EAA)	12	(A, EAA)	16	(EE, AA)
4	$(\cdot, EAAA)$	9	(E, AAA)	13	(A, AAA)	17	(EA, EA)
5	$(\cdot, AAAA)$					18	(EA, AA)
						19	(AA, AA)

For $m = 0.1, \epsilon = 0.001, f = 10, c = 1, b = 2$ and $k = 0.25$ we compute for this example $\mathbb{E}\Delta p = [MERK]^T v - v$. From the table below we can easily find out that

---

[6]Computations with all transitions are available upon request from the corresponding author, together with the computation of the invariant distribution that shows the difference in negligible.

this is not a linear function of $\sum p_i \left(1 - p_i\right) \frac{n_i}{N}$ and $\sum \left(p_i - \overline{p}\right)^2 \frac{n_i}{N}$; for states with identical values of these variables, the expected change in frequency differs.

	$\sum p_i \left(1 - p_i\right) \frac{n_i}{N}$	$\sum \left(p_i - \overline{p}\right)^2 \frac{n_i}{N}$	$\overline{p}$	$\mathbb{E}\Delta p$
$(\cdot, EEEE)$			0	0.001
$(\cdot, EEEA)$			0.25	-0.02993
$(\cdot, EEAA)$			0.5	-0.0396
$(\cdot, EAAA)$			0.75	-0.0295
$(\cdot, AAAA)$			1	-0.001
(E, EEE)	0	0	0	0.001
(E, EEA)	$\frac{1}{6}$	$\frac{1}{48}$	0.25	-0.0109
(E, EAA)	$\frac{1}{6}$	$\frac{1}{12}$	0.5	0.01039
(E, AAA)	0	$\frac{3}{16}$	0.75	0.06187
(A, EEE)	0	$\frac{3}{16}$	0.25	-0.06188
(A, EEA)	$\frac{1}{6}$	$\frac{1}{12}$	0.5	-0.07378
(A, EAA)	$\frac{1}{6}$	$\frac{1}{48}$	0.75	-0.05249
(A, AAA)	0	0	1	-0.001
(EE, EE)	0	0	0	0.001
(EE, EA)	$\frac{1}{8}$	$\frac{1}{16}$	0.25	0.01832
(EE, AA)	0	$\frac{1}{4}$	0.5	0.12476
(EA, EA)	$\frac{1}{4}$	0	0.5	-0.07129
(EA, AA)	$\frac{1}{8}$	$\frac{1}{16}$	0.75	0.03514
(AA, AA)	0	0	1	-0.001

The matrices M, E, R and K needed to calculate $\mathbb{E}\Delta p$ are very large. They can be obtained from the corresponding author, together with the computation of the invariant distribution and $\mathbb{E}\Delta p$ for every state.