Journal of Theoretical Biology 262 (2010) 391–397

Contents lists available at ScienceDirect

Journal of Theoretical Biology

journal homepage: www.elsevier.com/locate/yjtbi

Relatedness with different interaction configurations

Peter D. Taylor [a,*], A. Grafen [b]

[a] Department of Mathematics and Statistics, Queen's University, Kingston, Canada ON K7L 3N6
[b] Zoology Department, Oxford University, Oxford OX1 3PS, UK

ARTICLE INFO

Article history:
Received 10 November 2008
Received in revised form
8 October 2009
Accepted 8 October 2009
Available online 13 October 2009

Keywords:
Inclusive fitness
Hamilton's rule
Island model
Tag-based cooperation

ABSTRACT

In an inclusive fitness model of social behaviour, a key concept is that of the relatedness between two interactants. This is typically calculated with reference to a "focal" actor taken to be representative of all actors, but when there are different interaction configurations, relatedness must be constructed as an average over all such configurations. We provide an example of such a calculation in an island model with local reproduction but global mortality, leading to variable island size and hence variable numbers of individual interactions. We find that the analysis of this example significantly sharpens our understanding of relatedness. As an application, we obtain a version of Hamilton's rule for a tag-based model of altruism in a randomly mixed population. For large populations, the selective advantage of altruism is enhanced by low (but not too low) tag mutation rates and large numbers of tags. For moderate population sizes and moderate numbers of tags, we find a window of tag mutation rates with critical benefit/cost ratios of between 1 and 3.

1. Introduction

We have two objectives in this paper. The first concerns the construction of inclusive fitness models in a structured population in which individuals in different situations have interaction windows of different sizes and therefore have different numbers of interactions. In our investigation of this (in Sections 1 and 2), we look carefully at and gain a better understanding of the concept of relatedness in a "heterogeneous" population structure—one that lacks the internal symmetry among candidate actors that is so often implicitly assumed in theoretical inclusive fitness studies.

The second concerns the evolution of cooperative behaviour in a tag-based model in which individuals may have different tags and interact only with those who have the same tag. In case interactions are at random in the population as a whole, this will clearly lead to a situation in which individuals with different tags will have different numbers of interactions, and our analysis of the first problem will provide a solution to the second (Section 3). We will end with a general discussion of tag-based models.

Both of these objectives are of considerable interest. Ever since Hamilton's (1964) remarkable formulation of the inclusive fitness effect, its methodology has been much studied, both for its computational power and for the conceptual insights one can gain from its formulation. But the theory itself is sophisticated and any new example is an occasion for close study and deeper understanding of the method itself. The example studied here yields such understanding.

Much recent work (see the Discussion) has studied tag-based interactions, in which a behavioural trait, determined at one locus, is also affected by the recognition of a "tag" (e.g. skin colour, height) determined at another, or even culturally determined. There is evidence that individuals can recognize similar phenotypes and condition their behaviour accordingly (Matteo and Johnston, 2000; Lize et al., 2006; Sinervo et al., 2006, Sigmund, 2009) and there is interest in the general implications of this capacity on cooperative behaviour.

2. A general inclusive fitness formulation in a heterogeneous population structure

We work with a finite haploid population held at constant size with N individuals. Reproduction and mortality occur at a fixed rate. Time is divided into interaction intervals of length $1/N$. At the end of each such interval there is one birth and one death. The individual who gives birth is chosen according to relative fecundity and this is determined by a number of interactions which occur during that time interval. The population has constant size so every birth is met by a death and the individual who dies is chosen according to the population structure. In the model we will study here (Sections 2 and 3) this mortality will be random (global population regulation).

* Corresponding author. Tel.: +1 613 533 2434; fax: +1 613 533 2964.
E-mail addresses: peter.taylor@queensu.ca (P.D. Taylor), alan.grafen@sjc.ox.ac.uk (A. Grafen).

0022-5193/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jtbi.2009.10.012

We consider a behavioural locus at which there are two alleles, a resident allele B and a deviant allele A. The behaviour we wish to study is exhibited by an A-individual. It has a direct effect on the fecundity of certain individuals, and an indirect effect, since population size is constant, on the mortality of others. These effects are assumed to be small.

Individual fitness is calculated as the difference between fecundity and mortality. That is, individual fitness is the genetic representation of the individual in the population one time unit later, either through offspring or (negatively) through death. An offspring inherits the parental strategy except for a small bi-allelic mutation rate μ, small enough that we can ignore quadratic terms in μ.

The inclusive fitness formulation: Our objective is to track the changing frequency of the allele A under the action of selection. It seems clear enough that, in principle, the information required to do this is a comparison of the average fitness of A and B individuals and indeed a wonderfully compact formulation of Price (1970) for allele frequency change does exactly that. Hamilton's (1964) inclusive fitness method is closely related to Price's formula and provides an elegant and powerful heuristic approach. The formulation we work with here assumes there is no class structure in the population, that is, we do not have different sexes or ages or habitat qualities. Such heterogeneity requires us to account for variations in reproductive value (Taylor, 2009).

Take a focal A-individual (the *actor*) and add up the effects of its A-behaviour on all individuals in the population (the *recipients*), weighting each effect by the "relatedness" R of the actor to the recipient. What we get is called the inclusive fitness effect of the A-behaviour:

$$W_{\mathrm{I}} = \sum_j \Delta w_j R_j \tag{1}$$

where the sum is over all recipient's j, and Δw_j is the fitness effect on j, and R_j is the focal relatedness to j. The fundamental result of the general theory says that the sign of W_{I} will tell us whether selection will cause an increase (+) or decrease (−) in the frequency of A (Rousset and Billiard, 2000; Taylor et al., 2007a). This result requires a number of significant assumptions, notably that fitness effects are small (so that the result is valid to first order in these effects) and additive (between individuals—so that if my fitness is affected by the behaviour of two individuals the net effect is the sum of the two separate effects).

The key concept here is that of relatedness and some care is needed in formulating it. What is needed, of course, is a definition of relatedness which provides the fundamental result mentioned above, that the sign of W_{I} matches the sign of the selective change in the frequency of A. A standard approach for finite populations (Rousset and Billiard, 2000; Taylor et al., 2007a) arrives at the formula

$$R_j = \frac{G_j - \overline{G}}{1 - \overline{G}} \tag{2}$$

where G_j is the *coefficient of consanguinity* (CC) between the actor and recipient j defined as the probability that their two genes at the A–B locus are identical by descent (that is, have a common ancestor) and $\overline{G}$ is the average CC between the actor and the population—including the actor itself (Taylor et al., 2007a). Notice that this formulation of relatedness incorporates two normalizations, first the average relatedness of the actor to the population is zero and secondly (through the denominator $1-\overline{G}$) the relatedness of the actor to itself is 1. A consequence of the first is that the actor will, on average, have negative relatedness to a good fraction of the population. Thus relatedness measures not genetic identity *per se* but *relative* genetic identity.

Of course, in this formulation the focal actor is assumed to be "typical" in the sense that it is representative of all such actors. In practice there will certainly be variations among actors, both in regard to the fitness effects Δw_j and to the relatedness R_j to the recipient, and it is understood that both of these represent "averages" over all actors. The point we wish to make is that there are some subtleties and assumptions behind the taking of this average which require care. In order to discuss these, it is useful to have a simple example in mind. We consider the classic case of altruistic behaviour in which the actor confers a fecundity benefit b upon a recipient at personal fecundity cost c. Because population size is constant, these fecundity changes in actor and recipient will produce changes in mortality elsewhere. We can model this with four j-terms in the summation (1), actor and recipient and the two individuals with altered mortality effects.

The issue we are interested in here has to do, broadly speaking, with the relationship between the fitness effects and the relatedness. First of all, we remark that the average (over actors) of the product will not be the product of the averages unless the two terms are uncorrelated. For example, if an individual was somehow able to give a higher benefit b when it was more closely related to the recipient (see references cited above), this would have to be accounted for in formula (1). The model we study in Section 2 provides an interesting variant of this, in that different individuals interact different number of times. Of course we can always expect natural variation in numbers of interactions, but in the model we study here, this number happens to be positively correlated with the relatedness between the interactants. Individuals who tend to have more interactions also tend to be more closely related to their partners. How are we to handle this?

The critical insight is to distinguish between actors and individuals. An actor is an individual engaging in an interaction and the correct sample space, to track the effects of selection, is the space of actors, or, if you like, the space of interactions. For example, if we consider two individuals, I and J, and I interacts three times and J interacts once, then there are four actors, I counted three times and J counted once. With the set of *actors* as sample space, there is no longer a correlation between Δw_j and R_j and Eq. (1) is valid.

In this context we need to clarify the definition (2) of relatedness R_j. In the model of Section 2, different individuals not only have different relatedness to their partners, but also have different *average* relatedness to the population. Thus $\overline{G}$ in Eq. (2) is not the average CC to the population taken over all individuals, but is the average CC to the population taken over all actors.

So far, for our altruism example, Eq. (1) has four terms:

$$W_{\mathrm{I}} = bR_1 - cR_0 - (b\tilde{R}_1 - c\tilde{R}_0) \tag{3}$$

where R_1 and R_0 (=1) are the average relatedness of a random actor to its partner and to itself, and $\tilde{R}_1$ and $\tilde{R}_0$ are the average relatedness to the individual displaced by the offspring of the partner and of the actor itself. A special case arises if mortality is random and that is in fact the situation in Sections 2 and 3. In that case, $\tilde{R}_1$ and $\tilde{R}_0$ are zero, and the last two terms can be omitted leaving us with

$$W_{\mathrm{I}} = bR_1 - c \tag{4}$$

In many models, the assumption is often made, but not stated, that the mortality arising from the effects of the interaction is random. In Eq. (4) we get a positive inclusive fitness effect when

$$bR_1 > c \tag{5}$$

and this is Hamilton's classic 1964 rule for a selective increase in allele frequency.

3. An island model with variable patch size and two interaction protocols

Here we apply the analysis above to a population of constant size N structured as an island model with n patches of variable size (Fig. 1). An offspring either remains on its natal patch (with probability s) or disperses (with probability $d=1-s$) to a different patch chosen at random, but in neither case does it displace an existing breeder on that patch. Rather population regulation is global and at each birth an individual is chosen at random in the population to die. It is this dynamic that is responsible for the variation in patch size; however the coupling of birth and death events leads to a constant population size and a constant mean patch size of $m=N/n$.

This variable patch-size island structure has been investigated by Killingback et al. (2006) and Grafen (2007) in a model of a public goods game. In these applications, the patch structure was interpreted spatially as a demographic structure, and in this case one might feel the assumption of global mortality to be somewhat unrealistic, as it appears to require an unlimited elasticity of patch size—that is, if patches are physical entities, there will typically be a limited area in space they can occupy. However in the application to tag-based models we will investigate in Section 3, the structure is not spatial, rather we have a random-mixing population in which the patches are virtual collections of individuals with similar tags, and in this case the assumption of global mortality is completely natural.

Here we study the altruistic b–c interaction described above and consider two interaction protocols.

Protocol I: In each time step, each individual interacts with one partner chosen at random from within its patch.

Protocol II: In each time step, each individual interacts with *every* patchmate with some fixed probability θ, so that the expected number of interactions of an individual is proportional to its number of patchmates.

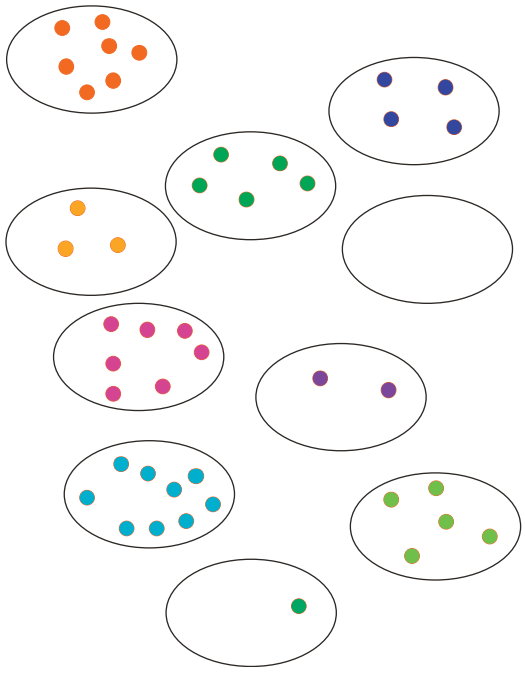

Fig. 1. *Island model with variable patch size*. In our application to tag-based cooperation in Section 2, this island structure exists in phenotype space, the patches containing all individuals with a fixed tag. For the purpose of interaction, the population mixes randomly, although an A-individual will behave altruistically only when encountering a partner with the same tag.

If the patches are all the same size these protocols give the same result (with additive fitness effects), but otherwise they do not. Protocol I is the model considered by Killingback et al. (2006) and Grafen (2007). Protocol II would arise if individuals *encountered* partners at random in the whole population but *interacted* only with those who were patchmates; in this case the average number of interactive encounters for an individual would be proportional to its number of patchmates. This is the situation we work with here. For both protocols, Hamilton's rule (5) applies and the condition for a positive W_I is

$$\text{Protocol I}: \quad R^{\mathrm{I}}b > c \tag{6}$$

$$\text{Protocol II}: \quad R^{\mathrm{II}}b > c \tag{7}$$

with the relatedness given in Eq. (2).

To display the difference between the two protocols, think about the relative contribution made by a patch of size k to the set of interactions. In protocol I the contribution is proportional to k and in protocol II it is proportional to $k(1-k)$. We get a simple and transparent version of the average over all instances of the behaviour if we change notation and let G_k be the average CC between patchmates in a patch of size k where the average is taken over all appearances of a patch of size k, and let $\overline{G}_k$ be the corresponding average CC of a patch k individual to the members of its population. If we finally let p_k be the frequency of patches of size k, the relatedness coefficients are

$$\text{Protocol I}: \quad G^{\mathrm{I}}=\frac{\sum_{k=2}^{N} kp_kG_k}{\sum_{k=2}^{N} kp_k} \quad \overline{G}^{\mathrm{I}}=\frac{\sum_{k=2}^{N} kp_k\overline{G}_k}{\sum_{k=2}^{N} kp_k} \quad R^{\mathrm{I}}=\frac{G^{\mathrm{I}}-\overline{G}^{\mathrm{I}}}{1-\overline{G}^{\mathrm{I}}} \tag{8}$$

$$\text{Protocol II}: \quad G^{\mathrm{II}}=\frac{\sum_{k=2}^{N} k(k-1)p_kG_k}{\sum_{k=2}^{N} k(k-1)p_k} \quad \overline{G}^{\mathrm{II}}=\frac{\sum_{k=2}^{N} k(k-1)p_k\overline{G}_k}{\sum_{k=2}^{N} k(k-1)p_k}$$

$$R^{\mathrm{II}}=\frac{G^{\mathrm{II}}-\overline{G}^{\mathrm{II}}}{1-\overline{G}^{\mathrm{II}}} \tag{9}$$

Computations: Eqs. (8) and (9) give us transparent expressions for the inclusive fitness effect under the two protocols, but they do not give us a feasible computational scheme for the relatedness coefficients except possibly for very small population size N. The standard way to calculate relatedness in a structured population is through recursive equations, but the variation in patch sizes in this population is a severe complication and feasible relatedness calculations for both protocols are not so easy to find.

Consider, for example, protocol I. In a recursive calculation of R^{I} we would at some point have two individuals in the same patch and need to know where they came from one time step earlier. Perhaps one of them is a new offspring, but to know whether it was an offspring of a parent in the same patch we need to know the size of the patch, so right away we have to cope with patch-size dependent relatedness coefficients. And then we run into the problem that the set of patch sizes is constantly changing. For this protocol, other methods seem to be required; as an example, Grafen (2007) has performed the R^{I} calculation (but with non-overlapping generations) using probability generating functions.

For protocol II, a modification of the standard recursive calculation was found by Antal et al. (2009) who implemented it in a finite one-dimensional "stepping-stone" population in which patches can house a variable number of individuals. In the supplementary materials A we use the Antal et al. approach to provide a recursive calculation of R^{II} in Eq. (9). The key idea is that the recursive approach can be made to work if the equations are

based, not on individuals, but on *pairs* of individuals. That is, in forming the recursive equations, instead of asking for the ancestor of each individual one time step ago, we ask for the ancestral *pair* of each *pair* of individuals one time step ago. The reason this gives us the protocol II result of course is that, as pointed out in Section 1, relatedness is calculated as an average over *interactions*, rather than over *individuals*, and in protocol II the number of interactions in a patch is proportional, not to the number of individuals in the patch but to the number of *pairs* of individuals in the patch.

To describe the approach in slightly more detail, we take a pair of individuals in the same patch or in different patches and ask whether the ancestral pair belonged to the same patch or to different patches. We let G_S and G_D be the average CC in each case (where this average is taken over all patch sizes in all population states) and we construct a pair of recursive equations for these. We do the same thing for the corresponding CC's $\overline{G}_S$ and $\overline{G}_D$ of an individual member of a random pair (in the same or different patch) to the population. Then if we are working with within-patch interactions (as we do in Section 3), we use the coefficients G_S and $\overline{G}_S$ in the relatedness expression (2); if we are working with between patch interactions, we use the coefficients G_D and $\overline{G}_D$.

It is important to point out that the recursive equations are formulated in a neutral population (with b and c equal to zero) it being generally infeasible to take account of the effects of positive b and c on the patch-size distribution. For small b and c (assuming smooth behaviour of the fitness functions) this will give us correct first-order results. Some results of the relatedness calculations are displayed Fig. 2 and are confirmed by population simulations (described in supplementary materials C) in Fig. 3.

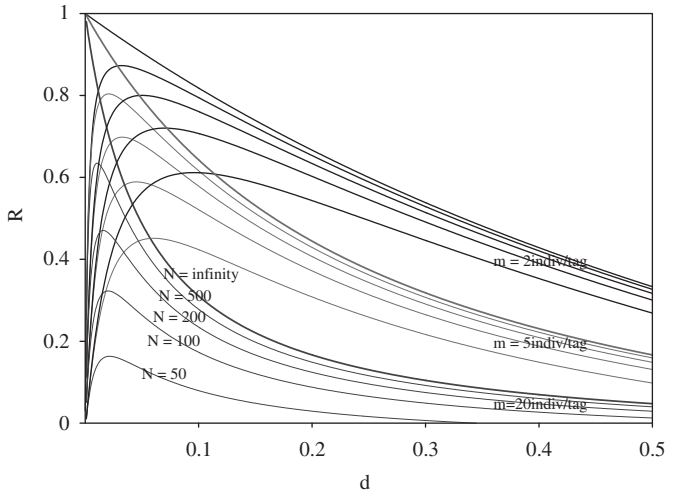

Fig. 2. Dependence of within patch relatedness R ($=R^{II}$ for protocol II) on offspring dispersal rate d in an island model (N individuals distributed on n patches) with global mortality and (therefore) variable patch size with average patch size $m=N/n$. Relatedness calculations are performed using the recursive equations developed in the supplementary materials A. Three sets of graphs are displayed for increasing values of m, with N increasing within each set. The graphs for increasing N tend to a limiting $N=\infty$ graph (Eq. (11)). Relatedness is seen to generally increase with population size N and with average patch size m. For fixed values of N and m, the dependence of R on d exhibits an interesting pattern. Except for very small d, relatedness increases with decreasing d, attains a maximum, and then sharply declines as d approaches zero. This decline is the result of a tendency, for very small dispersal rates, for the population to often find itself concentrated in a single patch in which (since average relatedness is zero) R will be negative. This decline is not seen in an infinite population in which relatedness tends to 1 as d approaches zero.

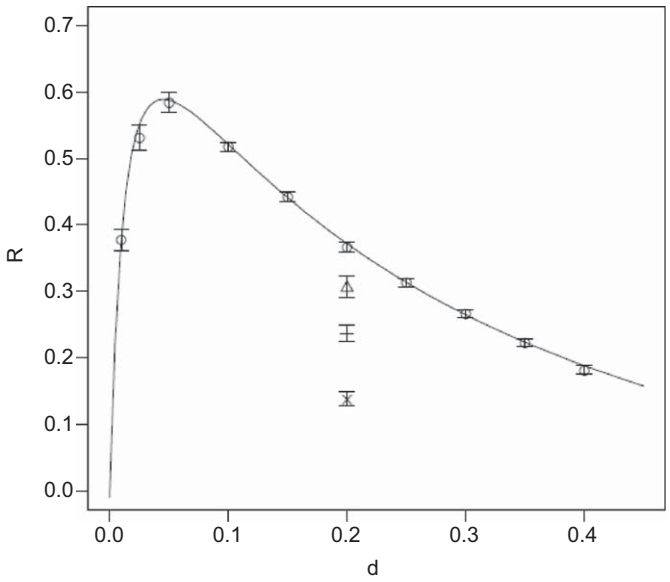

Fig. 3. Critical cost-benefit ratios (circles) obtained by computer simulation (supplementary materials C) plotted for the case of a population of size 100 with 20 available tags ($m=5$ individuals per tag) superimposed on the theoretical relatedness graph from Fig. 2. The height of the circles is the c/b ratio at which the average frequency of the alleles A and B equilibrates. The 95% confidence intervals are shown with a cross. The hypothesis that the true value lies outside the interval between the crosses can be rejected at the 5% significance level. These error bars of course depend on the number of simulations—if we did more the bars would get shorter. The present figure convinces us that the simulations get close to the true theoretical curve. The four points at $d=0.2$ belong to the recombination simulations. In each case the total probability that the offspring carries a different tag from the parent is 0.2. For each point the (tag-mutation, recombination) probabilities are: circle: (0.2, 0.0); triangle: (0.15, 0.05); bar: (0.10, 0.10); x: (0.05, 0.15).

4. Application to tag-based cooperation

We consider here a randomly mixed population in which each individual has, in addition to its strategy, a tag which belongs to a given set of n available tags. Offspring carry the same strategy and tag of the parent unless either of these mutates. We assume a strategy mutation rate μ, small enough that we can ignore quadratic terms in μ, and an independent but larger mutation rate d for tags. That is, an offspring has the same tag as its parent with probably $s=1-d$ and mutates to a different tag chosen at random with probability d.

Note that in the previous inclusive fitness analysis, we regarded the patch structure as spatial but in this application to tag-based cooperative behaviour, the population is randomly mixed and the island structure is a tag structure—patches correspond to sets of individuals with a fixed tag. In the terminology of Antal et al. (2009), the patch structure lives in phenotype space. As we have pointed out, with this set-up, random mortality is reasonable and expected.

Hamilton's rule (Eq. (7)) gives us a critical cost-benefit ratio of

$$\frac{c}{b}\Big|_{crit} = R^{II} \tag{10}$$

where for this model, we use protocol II as the proportion of an individual's encounters which result in interactions is proportional to the size of its tag set. Some results of the relatedness calculations are displayed in Fig. 2. The significant trends appear to be the dependence of relatedness on tag mutation rate d and on average "tag size", the number m of individuals per tag.

Dependence on tag mutation rate d: For the range of values we consider (of tag densities in the range $m=2$–20 individuals per

tag), the graphs show a maximum relatedness (maximum critical c/b ratio) at an intermediate mutation rate d, roughly in the region $0.002 < d < 0.01$. This can be explained as follows. For very low tag mutation rates, the tag sets behave like independent subpopulations with independent reproduction and global regulation. In this case, drift will cause one tag to dominate and the others to go extinct, so that much of the time the population is concentrated at a fixed tag. In that case, since average relatedness is 0, and the relatedness of an individual to itself is 1, the within-tag relatedness must be close to $-1/(N-1)$ and in particular will be negative, so that no benefit, no matter how large, can compensate for a positive cost. On the other hand, for high tag mutation rates, tags will be invaded too often by immigrants and within-tag relatedness does not get the chance to build to a high value.

Dependence on average tag size m: The graphs show that the more tags there are, the easier it is for cooperation to succeed. When there is one tag for every two individuals, cooperation is much more advantageous than when there is one tag for every twenty individuals. The main issue here is the probability that a tag-mutation will provide a tag that is currently unused, thereby allowing the mutant to establish a clone. If tag sets are frequently all occupied, there is much less chance for a disperser to found a new patch and therefore there is less chance for within-tag relatedness to build to high levels. Having said that, it must be noted that small values of m (no. of individuals per tag) will imply that individuals, on average rarely encounter partners with the same tag, so that selection will be very weak. However, even very small selective forces can sustain cooperative behaviour on those rare occasions of favourable encounters.

For large population size N and large numbers of tags n (but with $m=N/n$ fixed), the solutions to the protocol II recursions simplify to give

$$G \approx 1 - 2\mu N \frac{md}{1+d(m-1)}$$

$$\overline{G} \approx 1 - 2\mu N$$

giving the relatedness

$$R^{\mathrm{II}} = \frac{G-\overline{G}}{1-\overline{G}} \approx \frac{1-d}{1+d(m-1)} \tag{11}$$

to first order in μ. Note that it makes sense that if $d=1$, R^{II} must equal 0, and if $d=0$, R^{II} must equal 1. In fact, in an infinite population model, a rather simpler analysis is available (see supplementary materials B) and gives the above value of R^{II} directly. In Fig. 2 Eq. (11) appears for each value of m as a limiting curve (for large N). Eq. (11) also gives us the critical cost:benefit ratio for A to be selectively favoured, and in this it captures two of the trends discussed above, that small values of m (lots of tags) and small values of d (low tag-mutation rates) promote cooperation. What Eq. (11) does not capture is the selection against cooperation found in finite populations for very small values of d. That makes sense—in an infinite population, even very large tags will still be a vanishingly small proportion of the whole population and will have high relatedness because of the extreme population viscosity.

5. Discussion

A central puzzle in the evolutionary study of behaviour is the existence of altruism, or more precisely, the search for conditions under which altruistic behaviour is adaptive, and the understanding of how widespread such conditions might be in nature. The point is that a trait which caused its bearer to sacrifice personal fitness in order to increase the fitness of another ought to decline in frequency under the action of selection, so that the surviving traits will be effectively selfish.

A fundamental solution to this dilemma was offered by Hamilton (1964). If the altruist and recipient are genetically related so that the recipient has a high enough probability of also bearing the altruistic allele, then the altruistic action can, on average, lead to an increase in the frequency of the allele. This scenario, of altruism and cooperation between relatives, is played out in many different ways in natural populations; Hamilton's genius, and his enduring legacy, was to understand that such interactions can be formulated as maximizing principles, as if the actor chooses its behaviour to maximize its inclusive fitness—a relatedness-weighted average of its effects on all individuals (Grafen, 2006). In the simple altruism model discussed here, the inclusive fitness effect can be written in the form $W_1=Rb-c$, and from this we get "Hamilton's rule" which says that the actor will choose altruism when $Rb > c$, when the benefit b weighted by the relatedness R between interactants exceeds the cost c to the altruist.

But the idea raises many problems, a central one of which is the question of how I am to recognize my close relatives in order to restrict my beneficence to them? There have been two types of answers given to this, both discussed by Hamilton—one through proximity, the other through similarity—but both raise further difficulties.

Interacting with neighbours: The first of these, proximity, argues that in a "viscous" population in which individuals over their lifetime do not move far from their birth-place, one's neighbours are often one's relatives, so one can simply interact with those who live nearby. In this case, if the b/c ratio is sufficiently high it will exceed the threshold $1/R$ needed for altruism to be favourably selected, where R is now average relatedness to ones neighbours. The difficulty raised with this is that in a viscous population, neighbours tend also to be competitors in that their offspring compete with one another for resources (food and space) and a benefit given to a neighbour which results in additional offspring will provide greater competition for the altruist and possibly undo the positive effect on the altruistic allele (Hamilton, 1971). In fact there is a general result which shows that in a population with a "transitive" structure of breeding sites, these two opposing forces cancel one another out (Wilson et al., 1992; Taylor, 1992a, b; Taylor et al., 2007b; Grafen and Archetti, 2008), so that interaction with neighbours can support altruistic behaviour only under special circumstances. [More precisely what "transitive" needs to mean for the above result is that for any pair of breeding sites (i, j), there is a bijective transformation of the space of sites mapping i to j and preserving offspring movements and interaction probabilities.]

Interacting with those who are similar: The second way to "recognize" relatives works by the principle that close relatives tend to share a number of recognizable characteristics. Turning this around, interactants that share a number of characteristics might well be relatives, and in particular, might well have the same cooperative tendencies. Introduced as the "armpit effect", (Dawkins, 1982) this study now goes under the name of tag-based models of cooperation. In such models, each individual has a tag, possibly multi-dimensional, and interacts only with those who have the same or nearby tag. The difficulty here is that whatever the mechanism that links behaviour and tag, whether genetic or otherwise, there are apt to be ways of cheating, of having the right tag but not the right behaviour. Much recent work has been devoted to the study of the dynamics between these two traits to clarify the circumstances under which tag-based models can promote cooperative behaviour. The first formal model (Riolo et al., 2001) assumed a randomly mixed population, a one-dimensional continuum of tags (similar to Antal et al., 2009) and

allowed the threshold distance (for cooperation) to evolve. However their model excluded the possibility of cheating—everyone was a cooperator when they encountered a partner with a close enough tag. In this set-up, the tag and the cooperative behaviour come hand-in-hand; such mechanisms, in which tag and behaviour emerge from the same pleiotropic gene, are referred to as "green beard" models. Green beard genes are hard to find, but see Keller and Ross (1998) for a classic example in the red fire ant. More recent models have found random mixing to be rather hostile to tag-based cooperation but have generated a selective advantage for cooperation through the use of population structure (Crozier, 1986; Grafen, 1990; Axelrod et al., 2004; Jansen and van Baalen, 2006; Hammond and Axelrod, 2006; Rousset and Roze, 2007; Gardner and West, 2007; Traulsen and Nowak, 2007).

Our model works with a randomly mixed population in which individuals interact only when they encounter partners with the same tag. We obtain a window of parameter values within which cooperation is selectively favoured for a range of b/c values between, say, 1 and 3, but this window is often quite small. Generally we find that we need large numbers of available tags and a low (but not too low) tag mutation rate. Two recent models (Antal et al., 2009; Tarnita et al., 2009) offer different extensions which are in ways more natural or realistic. Antal et al. (2009) puts a one- or two-dimensional structure on the set of tags so that the "patches" are arrayed on a grid in "phenotype space" and tag-mutation is to neighbouring tags. This complicates the calculation of relatedness but the results that emerge are simple and elegant. They provide a complete analysis of the one-dimensional case and obtain the result that for large populations cooperation is selectively advantageous when $b/c > 1+2/\sqrt{3} = 2.15$. One of the remarkable things about their model is that the parameter m is missing! The essential reason for this is that the organic nature of the population dynamics allows the parameters to interact in a natural way. The one-dimensional grid structure occupies an infinite line and while the population is allowed to drift along the line, it stays together in a cluster of mean width $\sqrt{dN}$ where tag-mutation d is restricted to the two immediate neighbours, and this automatically fixes the average number of individuals per *occupied* tag. Tarnita et al. (2009) extend the model in a different way. Like us, they have a discrete unconnected tag set but they allow individuals to share multiple tags, and interactions occur only if "enough" tags are shared.

It is instructive to relate the general equilibrium condition (70) in Antal et al. (2009), to our general relatedness formula (2). The condition of Antal et al. (2009) can be written (divide top and bottom of (70) by z) as

$$\left(\frac{b}{c}\right)^{*} = \frac{(N-2)(1-\hat{G})+(1-G)}{(N-2)(G-\hat{G})-(1-G)} \tag{12}$$

where G is the CC between two interactants and $\hat{G}$ is the average CC's of the focal to the rest of the population *excluding both the focal and its interactant* (noting that the individuals k, q and l in their Eq. (1) are distinct). According to Hamilton's (1964) rule, the right side of condition (12) should equal $1/R$ where R is the relatedness between interactants. To see that this is so, let $\overline{G}$ be the average CC of the focal interactant to the entire population. Then $N\overline{G} = (N-2)\hat{G}+1+G$, and if this is used to eliminate $\hat{G}$ in Eq. (12), we get

$$\left(\frac{b}{c}\right)^{*} = \frac{1-\overline{G}}{G-\overline{G}} \tag{13}$$

According to our general expression (2), the right-hand-side of Eq. (13) is $1/R$, as expected.

The CC calculations: We work with two protocols. In protocol I, an individual interacts with a random patchmate; in protocol II, an individual interacts with each patchmate with some fixed probability. Each gives rise to a different average relatedness coefficient. In both cases, R is a weighted average of the relatedness between patchmates over all patches, but in protocol I, in which the number of interactions in a patch is proportional to the number of individuals in the patch, a patch of size k receives a relative weight of k, whereas in protocol II, in which the number of interactions in a patch is proportional to the number of *pairs* of individuals in the patch, it receives a relative weight of $k(k-1)$. These weightings will give different results when within-patch relatedness co-varies with patch size. The resulting formulae are nicely displayed in Eqs. (8) and (9), but these equations do not provide a feasible calculation of the average coefficients, and in both cases the calculation presents a challenge.

The effect of recombination: Recombination and tag-mutation both have the effect of giving the offspring a new tag, different from its parental tag (and thereby they both tend to break up any association between tag and strategy), but they work in different ways. Mutation chooses the new tag at random from a set of available tags, whereas recombination weights the choice of new tag by the number of individuals currently holding that tag. If tag-sets were roughly the same size, this would amount to the same thing, otherwise they are somewhat different. Because of this dependence of recombination on the current distribution among tags, it is not easy to incorporate recombination in an inclusive fitness analysis, but see Rousset and Roze (2007) for an analytic model of recombination in a patch-structured population. We have, however, run a few simulations (Fig. 3) which show recombination to be much more hostile to cooperation than is tag-mutation. One possible reason for this is that tag-mutation allows the offspring of an A-individual to obtain an "unoccupied" tag and thereby establish a pure A-colony. Recombination cannot do that.

The specification of relatedness: What qualifies R_j in Eq. (2) as a relatedness coefficient are two critical normalizations—first that average relatedness to the population is zero and secondly that relatedness to self is one. The significance of this is that an actor can pose the two questions: what difference would it make if I had chosen my partner at random or if I had chosen my partner to be myself? And the answer can be found by replacing R_j with 0 or 1. To be more precise, that answer actually only works *on average*. It applies to an individual engaged in a random interaction.

Having said that, it is worth pointing out that if an actor can recognize and respond to different actor configurations, it seems reasonable that different types of actor might evolve different "conditional" behaviours, and we could analyse the situation for each type separately. Even for the tag-based model discussed here, it is reasonable to suppose that an individual who tended to encounter many individuals with the same tag might begin to behave differently from one who encountered only a few.

Finite populations: We have pointed out that relatedness must be calculated as an average over all interactions. In finite populations, we need to make clear exactly how this average is calculated. In our patch structured population, if we let the *state* of the population be the number of patches of different sizes, then over time the state of the population will continually change, and it might well be (for a certain b and c) that the frequency of the allele A will increase from some states and decrease in others. A general theory for allele frequency change in finite populations (Rousset and Billiard, 2000; Taylor et al., 2007a) proposes that the overall direction of allele frequency change should be calculated as an average over all possible states of the rate of change at each state. What this tells us is that the average in Eqs. (1) and (2) needs to be over all interactions in any fixed state and then this in turn needs to be averaged over all population states, or equivalently over time. For example the average CC G_k between patchmates in a patch of size k

in Eqs. (8) and (9) must be taken over all appearances of a patch of size k over all time.

Reproductive value: The structure of breeding sites in our population is inhomogeneous (due to variation in patch size) and in this case we can typically expect different breeding sites to have different reproductive value (RV) and this must be accounted for in the inclusive fitness calculation (Taylor, 2009). For our particular model, however, the RV of a breeding site is independent of patch size, at least in the neutral population (with $b=c=0$). The reason for this is that fecundity is independent of patch size and since there are no local competitive effects, this is also the case for mortality.

Acknowledgements

This work was partially supported by the Programme in Evolutionary Dynamics at Harvard University, by the Leverhulme Foundation and St. John's College, Oxford, and by the Natural Sciences and Engineering Research Council of Canada.

Appendix A. Supplementary material

Supplementary data associated with this article can be found in the online version at 10.1016/j.jtbi.2009.10.012.

References

Antal, T., Ohtsuki, H., Wakeley, J., Taylor, P.D., Nowak, M.A., 2009. Evolutionary game dynamics in phenotype space. PNAS 106, 8597–8600.
Axelrod, R., Hammond, R.A., Grafen, A., 2004. Altruism via kin-selection strategies that rely on arbitrary tags with which they coevolve. Evolution 58, 1833–1838.
Crozier, R.H., 1986. Genetic clonal recognition abilities in marine invertebrates must be maintained by selection for something else. Evolution 40, 1100–1101.
Dawkins, R., 1982. The Extended Phenotype. W.H. Freeman, San Francisco, CA.
Gardner, A., West, S.A., 2007. The decline and fall of genetic kin recognition. Current Biology 17, 810–812.
Grafen, A., 1990. Do animals really recognize kin? Anim. Behav. 39, 42–54.
Grafen, A., 2006. Optimization of inclusive fitness. J. Theor. Biol. 238, 541–563.
Grafen, A., 2007. Detecting kin selection at work using inclusive fitness. Proc. R. Soc. B 274, 713–719.
Grafen, A., Archetti, M., 2008. Natural selection of altruism in inelastic homogeneous viscous populations. J. Theor. Biol. 252, 694–710.
Hamilton, W.D., 1964. The genetical evolution of social behaviour I. J. Theor. Biol. 7, 1–16.
Hamilton, W.D., 1971. Selection of selfish and altruistic behaviour in some extreme models. In: Eisenberg, J.F., Dillon, W.S. (Eds.), Man and Beast: Comparative Social Behaviour. Smithsonian Press, Washington, DC, pp. 57–91.
Hammond, R.A., Axelrod, R., 2006. Evolution of contingent altruism when cooperation is expensive. Theor. Popul. Biol. 69, 333–338.
Jansen, V.A.A., van Baalen, M., 2006. Altruism through beard chromodynamics. Nature 440, 663–666, doi:10.1038/nature04387.
Keller, L., Ross, K.G., 1998. Selfish genes: a green beard in the red fire ant. Nature 394, 573–575.
Killingback, T., Beri, J., Flatt, T., 2006. Evolution in group-structured populations can resolve the tragedy of the commons. Proc. R. Soc. B 273, 1477–1481.
Lize, A., et al., 2006. Kin discrimination and altruism in the larvae of a solitary insect. Proc. R. Soc. B 273, 2381–2386.
Matteo, J.M., Johnston, R.E., 2000. Proc. R. Soc. Lond. 267, 695–700.
Price, G.R., 1970. Selection and covariance. Nature 227, 520–521.
Riolo, R.L.., Cohen, M.D., Axelrod, R., 2001. Evolution of cooperation without reciprocity. Nature 414, 441–443.
Rousset, F., Billiard, S., 2000. A theoretical basis for measures of kin selection in subdivided populations: finite populations and localized dispersal. J. Evol. Biol. 13, 814–825.
Rousset, F., Roze, D., 2007. Constraints on the origin and maintenance of genetic kin recognition. Evolution 61 (10), 2320–2330.
Sinervo, B., et al., 2006. Self-recognition, color signals, and cycles of greenbeard mutualism and altruism. Proc. Natl. Acad. Sci. USA 103, 7372–7377.
Sigmund, K., 2009. Sympathy and similarity: the evolutionary dynamics of cooperation. Proc. Natl. Acad. Sci. USA 106, 8405–8406.
Tarnita, C.E., Antal, T., Ohtsuki, H., Nowak, M.A., 2009. Evolutionary dynamics in set structured populations. PNAS 106, 8601–8604.
Taylor, P.D., 1992. Altruism in viscous populations—an inclusive fitness model. Evol. Ecol. 6, 352–356.
Taylor, P.D., 1992. Inclusive fitness in a homogeneous environment. Proc. R. Soc. B 249, 299–302.
Taylor, P.D., Day, T., Wild, G., 2007a. From inclusive fitness to fixation probability in homogeneous structured populations. J. Theor. Biol. 249, 101–110, doi:10.1016/j.jtbi.2007.07.006.
Taylor, P.D., Day, T., Wild, G., 2007b. Evolution of cooperation in a finite homogeneous graph. Nature 447, 469–472.
Taylor, P.D., 2009. Decompositions of Price's formula in an inhomogeneous population. J. Evol. Biol. 22, 201–213, doi:10.1111/j.1420-9101.2008.01640.x.
Traulsen, A., Nowak, M.A., 2007. Chromodynamics of cooperation in finite populations. PLoS One 2, e270.
Wilson, D.S., Pollock, G.B., Dugatkin, L.A., 1992. Can altruism evolve in purely viscous populations? Evol. Ecol. 6, 331–341.