VICTORIA UNIVERSITY OF WELLINGTON
Te Whare Wananga o te Upoko o te Ika a Maui

Computer Science

PO Box 600
Wellington
New Zealand

Tel: +64 4 471 5328
Fax: +64 4 495 5232
Internet: office@mcs.vuw.ac.nz

The Evolution of Policing

Richard Mansfield

Supervisor: Marcus Frean

28 February 2001

Submitted in partial fulfilment of the requirements for
Bachelor of Science with Honours in Computer Science.

Abstract

The concept of policing is intended to help explain how cooperative groups have evolved from competing individuals. Recent theories have provided mathematical models of policing which propose that individuals invest resources in the repression of competition within their local group. These models are investigated in this project using agent-based computer simulations, and some alternative models of policing are proposed.

1 Introduction

Groups of self-interested individuals who are competing for a limited resource are often caught in a prisoner's dilemma, in which every individual in the group would be better off if the whole group were more cooperative, but any individual who becomes more cooperative will be punished with a smaller share of resources when others in the group are more competitive. In the absence of any other mechanism to enhance cooperation, selfish competitive behaviour will spread completely throughout the group, and individual and group fitness will stay low. In spite of this, there are many examples in nature of large cooperative groups, made up out of smaller evolutionary units, that look as if they would have had to have developed in a prisoner's dilemma-like scenario. Evolutionary theory needs to be able to explain how competition among the lower-level units could have been suppressed for long enough for cooperative groups to form. For example, evolutionary theory has so far not explained how small individual replicator molecules managed to group together into cells without the process being subverted by competitive free-riding replicators. Nor has it explained how competition between cells was suppressed to enable the creation of multicellular organisms, or how cooperative social groups made up of individual, selfish animals have come about. The concept of 'policing' is one possible answer.

The term 'policing' has been used in evolutionary biology to describe the behaviour of honey bee workers, who will sometimes attack other workers in the same colony that lay their own eggs rather than tending to the queen's eggs[6]. These 'policing' workers are acting for the benefit of the entire colony at a personal cost to themselves. Steven Frank's model of policing is the subject of this project. Frank develops an abstract theoretical model that he hopes will apply to examples of policing found in nature, such as honey bees. This project started out with the aim of using an agent-based simulation to probe Frank's model, but the simulation itself suggested new ways of formalising the concept of policing.

This section is a general introduction to Frank's model of policing. Section 2 describes the method that he uses to analyse the consequences of his models mathematically. In section 3, I describe the agent-based simulations that I performed on Frank's model, and in section 4, I suggest some alternative models of policing.

1.1 Policing

Frank proposes an abstract trait called policing as a possible mechanism to counter the costs of competition within a group [2],[3]. He describes a mathematical model in which every agent has two traits: a degree of competitiveness and a degree of investment in policing. He claims that his model provides the necessary conditions for larger cooperative groups to form. Policing is a trait that is supposed to suppress the negative effects of competition, but Frank is not concerned with exactly how this suppression works — he is only interested in finding out under what circumstances policing is favoured by evolution.

The idea is that an individual can choose to invest in policing at a personal cost to their own fitness. Frank's model assumes that policing delivers more group benefits, *as if* other members of the group had been less competitive, so the cost of this investment can be paid back to the policing individual through benefits derived from being part of a successful group. This model relies on the concept of group selection, which is the idea that groups can have different levels of fitness and that natural selection operates on these groups. Group selection was once widely disbelieved in evolutionary biology (see for example Dawkins' critique of Wynne-Edwards formulation of group selection in [1]), but it is now commonly accepted that natural selection can operate on a nested hierarchy of units such as groups, individual organisms and genes [10]. The differential fitness between groups, if the groups have certain properties, can have a strong effect on evolution, and the contribution of group selection to evolution can in fact be measured [5].

Investment in policing is an example where group selection may have a strong effect. The individual investing in policing pays a personal cost, and therefore lowers his fitness relative to the rest of his group. It might be thought that policing could never take off under these circumstances, because free-riding individuals who don't invest in policing will always be fitter and will eventually drive the policing-investers to extinction. But this will only be true if the negative effect on the fitness of the policing individual from the within-group selection actually outweighs the positive effects from the between-group selection. In other words, if groups that have policing are sufficiently better off than those that don't have policing, policing could actually take off in the population.

Investment in policing will be worthwhile when the competition within the group is high, because the fitness that an individual loses from their investment in policing will often be less than the potential gain in fitness from the suppression of the effects of high group competition. If for some reason group competitiveness is very low, the cost of policing may not be paid back in extra benefits.

1.2 Frank's simple model

Frank starts with a fitness equation for an individual in a group without policing:

$$w_{ig} = \frac{z_{ig}}{z_g}(1 - z_g) \tag{1}$$

where—

- w_{ig} is the fitness of the i^{th} individual in the g^{th} group;
- z_{ig} is i's level of competitiveness; and
- z_g is the average level of competition in the g^{th} group.

The $(1 - z_g)$ part represents the average group productivity, which declines as the group becomes more competitive internally. The z_{ig}/z_g represents i's share of the total group resources, which increases as i becomes more competitive relative to the group average. This model just describes a prisoner's dilemma, except that Frank states his model in terms of competitiveness, while the prisoner's dilemma is normally stated using 'cooperation'. In this equation it looks as if Frank intends the competitiveness, or z, trait to range only between 0 and 1, because of the $1 - z_g$ bracketed term. We can therefore think of an individual's level of cooperation as $1 - z_{ig}$, and the average level of cooperation in the group as $1 - z_g$.

1.3 Frank's policing model

The above model has been extended to include the effects of policing, so that

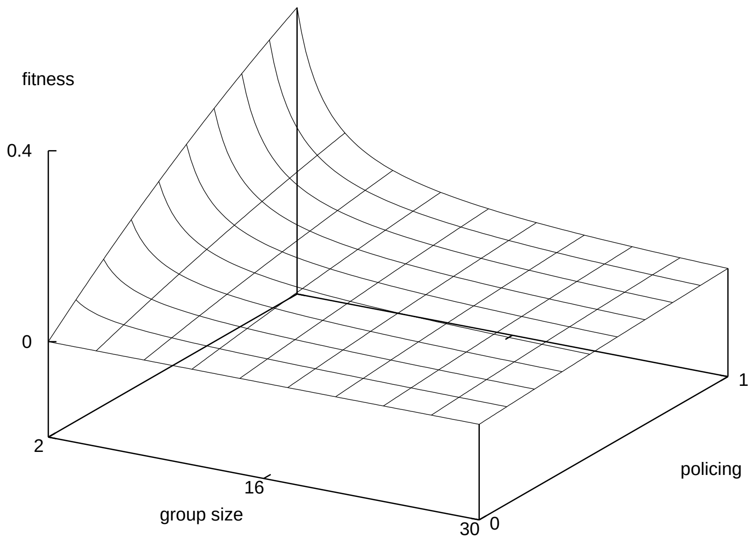

Figure 1: The fitness of a single policing individual in an uncooperative group when $k = 0.2$. Policing is more stronly favoured when the group size is small.

$$w_{ig} = [a_g - ka_{ig} + (1 - a_g)\frac{z_{ig}}{z_g}][1 - (1 - a_g)z_g] \tag{2}$$

where—

a_{ig} is i's investment in policing (where i is a member of the g^{th} group);
a_g is the average level of policing in the g^{th} group; and
k is the cost of policing.

The second set of terms in square brackets, like the second part of equation (1), is supposed to measure the success of the group g. The negative effects of z_g on the size of the group resources are now moderated by policing; only a proportion $(1 - a_g)$ of the group resources (the ones obtained from unpoliced interactions) are lost to group competition.

The terms in the first set of square brackets are supposed to be proportional to the individual's fitness within the group. The idea here is that there is a proportion of the group interactions, a_g, which are policed, and the resources gained from these interactions are shared fairly among all the members of the group. A further proportion $(1 - a_g)$ of the interactions are unpoliced, and the gains from these interactions are up for grabs by competing individuals (so these unpoliced interactions are multiplied by z_{ig}/z_g, whereas the a_g policed interactions are multiplied by 1 so that they that they are not subject to within-group competition). The $-ka_{ig}$ term is the cost that each individual pays for their own investment in policing.

1.4 Policing is favoured in a competitive population

Frank doesn't make it clear exactly how his model encourages policing, but we can see it by looking at how a single individual with the policing trait could prosper in a competitive population. Suppose i is the only member of the group who contributes anything towards policing, and the size of the group is n_g. If i's contribution to policing is a_{ig}, then the average level of policing in the group, a_g, will be equal to a_{ig}/n_g. In a totally uncooperative group, $z_g = z_{ig} = 1$. Plugging these into equation (2) we get

$$w_{ig} = (1 - ka_{ig})\frac{a_{ig}}{n_g} \tag{3}$$

If i had spent nothing on policing, his payoff would have been zero, so policing will be favoured for i whenever

$$(1 - ka_{ig})\frac{a_{ig}}{n_g} > 0$$

This is true for any k less than one, as long as a is less than one. In a small group, the immediate gain in fitness for the individual who invests in policing will be high, but in a large group, there will be little to gain for an individual who invests in policing, because the proceeds from his investment will be spread out among the whole group. Figure 1 shows the fitness of a lone policing individual for different levels of a_{ig} and various sizes of totally uncooperative groups. It shows that policing of the kind Frank describes should be much more strongly selected for in small groups.

1.5 Policing requires group selection

The single policing individual described above whose fitness is described by equation (3), will have lower fitness than the rest of his group, who will have fitnesses equal to a_{ig}/n_g. This is just because he is the only one paying the cost of policing ka_{ig}, while everyone else in the group is free riding; they are reaping the benefits from the policed interactions without paying for them. But the policeman's fitness is still higher than it would be if he were not policing at all, and higher when compared to any member of any non-policing competitive group that might be out there in the wider population. If the population size is restrained for some reason, (in the case of animals this could be because of finite available food or territory), and if the population does not live or interact in smaller groups, then policing could never take off, for the simple reason that a policing individual's fitness will always be worse relative to the rest of the population. Even if the population size is unbounded, the policing trait could only increase in terms of sheer numbers, but proportionally it would decrease. In order for policing to prosper as a proportion of the population, we need policing and non-policing groups to have different relative levels of success, and we need this difference to outweigh the relative difference between policing and nonpolicing individuals within the groups.

2 Finding optimal levels of policing and competition

This section describes a mathematical method [9] for finding the optimal level of policing and competitiveness from a fitness equation. This method is used to find optimal levels of policing and competition in any model, like Frank's, which is described by a fitness equation.

2.1 Relatedness

The theory of kin selection was developed to explain certain types of altruistic behaviour towards relatives [4]. In some situations it is worthwhile for an individual to put some investment into its relatives' reproductive success, because some of its genes are shared with those relatives. For example, because my sister probably shares half of my genes, investing in her reproductive success is about half as good for me (other things being equal) as investing in my own reproductive success. Or in other words, it should pay off for me to invest half as much in her reproductive success as I would in my own. More generally, the amount that A should be prepared to invest in B's reproductive success is proportional to the *relatedness* between A and B. In this case the relatedness, r, represents the chance that a gene is shared between A and B (beyond the average genetic relatedness for the species).

2.2 Group relatedness

In terms of groups, relatedness is just the chance that a gene in one member of the group will be shared by another random member of the group, or, the proportion of the group members that will have any given gene. Taylor and Frank claim that relatedness can also be thought of as the change in the level of a group characteristic (the average for all members of the group) with respect to an individual characteristic [9]. For example, consider the competitiveness trait, z, used by Frank in his fitness equations (1) and (2). The group relatedness is $\frac{\mathrm{d}z_g}{\mathrm{d}z_{ig}}$, the gradient of group phenotype on individual phenotype, because if we alter an underlying gene so that i's level of competitiveness z_{ig} changes, then that alteration will have an effect on z_g that can be determined by the number of copies of that gene which are spread throughout the rest of the group, which is equal to r. This will hold for any trait in the population. For the policing and competitiveness traits that Frank is interested in, $r = \frac{\mathrm{d}z_g}{\mathrm{d}z_{ig}} = \frac{\mathrm{d}a_g}{\mathrm{d}a_{ig}}$.

Any high level of group relatedness is usually thought of as being driven by an assortment mechanism such as *kin recognition*. This occurs when related individuals can recognise one another, and they prefer to get together into groups to perform cooperative tasks. Under these conditions r will be high.

2.3 Evolutionary stable strategies

When a population, whose fitness is determined by an equation like Frank's, reaches equilibrium, we have what Maynard Smith calls an evolutionary stable strategy (ESS) [7]. An ESS is an individual strategy that if the vast majority of the population followed it, could not be invaded by a small number of individuals following any different strategy. In Frank's equation (1) model, an individual's strategy can be described very simply; it is just a single number, z_{ig}. The ESS for this model is at $z = 1$, because in a large population of individuals following this strategy, any small number of individuals with a lower level of z will do badly and will be unable to invade the population. With an ESS model, we're entitled to assume that at equilibrium *all* individuals will have the same phenotype, or trait value. In Frank's models, this means that $z_g = z_{ig}$ and $a_g = a_{ig}$ at the equilibrium point.

2.4 The Taylor and Frank method for finding an evolutionary stable strategy

Taylor and Frank give a method for finding the optimum level of a trait in a population that has reached equilibrium [9]. The method is

(1) Find the partial derivative of fitness equation with respect to the trait concerned.

(2) Replace the 'phenotypic derivatives' of the trait, such as $\frac{\mathrm{d}z_g}{\mathrm{d}z_{ig}}$, with r.

(3) Remove the subscripts on all the trait variables (because at equilibrium $z_g = z_{ig}$, etc.)

(4) Find the equilibrium values by solving the derivative for the trait concerned at the point where the derivative is equal to zero.

2.5 Applying the Taylor and Frank method

To give a detailed example of the application of Taylor and Frank's method, I will show how it can be used to find the evolutionary stable level of a in the model specified by equation (2). We start by finding the partial derivative of the fitness w_{ig} with respect to the policing trait a_{ig}. Following step (1), we get

$$\frac{\partial w_{ig}}{\partial a_{ig}} = (a_g - ka_{ig} + (1-a_g)\frac{z_{ig}}{z_g})(z_g\frac{\mathrm{d}a_g}{\mathrm{d}a_{ig}}) + (\frac{\mathrm{d}a_g}{\mathrm{d}a_{ig}} - k - \frac{z_{ig}}{z_g}\frac{\mathrm{d}a_g}{\mathrm{d}a_{ig}})(1-(1-a_g)z_g)$$

Note that because a_g is not independent of a_{ig}, we are not allowed to treat a_g as a constant. Step (2) is replacing all occurrences of $\frac{\mathrm{d}a_g}{\mathrm{d}a_{ig}}$ by r, and gives us

$$\frac{\partial w_{ig}}{\partial a_{ig}} = (a_g - ka_{ig} + (1-a_g)\frac{z_{ig}}{z_g})z_g r + (r - k - \frac{z_{ig}}{z_g}r)(1-(1-a_g)z_g)$$

Step (3), losing all the subscripts (because at equilibrium $z_{ig} = z_g$ and $a_{ig} = a_g$) gives

$$\frac{\partial w}{\partial a} = zr(1-ka) - k(1-(1-a)z) \tag{4}$$

The gradient for z, which can be derived similarly, is

$$\frac{\partial w}{\partial z} = (1-a)[(1-(1-a)z)(1-r)/z - r(1-ka)] \tag{5}$$

	$r < 1 - k$	$r > 1 - k$
a	1	0
z	$(1-r)/[r(1-k)]$	$1-r$

Table 1: Optimal values for a and z in Frank's model of policing (equation (2)). Individuals with these values will be following an evolutionary stable strategy.

The evolutionary stable level of a and z will occur when the two gradients above are equal to zero[1]. Working out under exactly what conditions both are zero is not at all straightforward. From equation (5), assuming a is not exactly 1, it follows that $\frac{\partial w}{\partial z} = 0$ when

$$zr(1 - ka) = (1 - (1 - a)z)(1 - r)$$

and if we substitute this into equation (4) we get

$$\frac{\partial w}{\partial a} = (1 - (1 - a)z)(1 - r - k)$$

The first part is positive, which shows us that the sign of the gradient for a depends only on the values of r and k. The gradient for a is positive whenever $r < 1 - k$, and a will be driven up to 1. Whenever $r > 1 - k$, the gradient is negative and a will be driven down to 0. The equilibrium values for z can be found by looking at equation (5) for these two cases. If we consider the case when $r < 1 - k$, and substitute $a = 1$ back into equation (5), we get no gradient, but as a tends towards 1, z will tend towards $(1 - r)/[r(1 - k)]$. Similarly, if we substitute $a = 0$ into equation (5), we can see that whenever $r > 1 - k$, we get $(1 - z)(1 - r)/z = r$, and $z = 1 - r$. Table 1 summarises the expected equilibrium values for a and z at different values of r.

Frank's primary result is thus that policing is favoured whenever the group relatedness is low, and that policing is not favoured when relatedness is high. This is what we expect to find, because (according to kin selection theory) competition should be lower among relatives, reducing the need for policing. When relatedness is low, within-group competition will be high, the potential between-group benefits will be high, and it will be worthwhile paying the cost of policing.

2.6 Three further comments on Frank's policing model

One important thing to notice about Frank's model is that the traits a and z will not always correspond to what we think of in real life as policing and competition. In the equation (2) model, it cannot be said that policing actually suppresses competition. It alleviates the negative fitness effects of competition, but it actually increases competition by driving z higher than it would have been without policing (perhaps in the same way that crime becomes

[1] In Frank's model, $0 \leq a \leq 1$, so $\frac{\partial w}{\partial a} = 0$ may not in fact be reachable for any valid a. In this case a will saturate to 0 or 1 depending on the sign of the gradient. z, however, is not actually restricted to be less than 1 in Frank's model; see section 2.6.

more brutal during periods of prohibition when policing is overly heavy-handed). This is unrealistic for many types of policing, because we often find that a lot of the behaviour that we would like to call 'policing' really does suppress competition. Frank partially addresses this complaint with a further model in which competitiveness also has an associated cost (see section 3.10).

There are also a couple of discrepancies between Frank's model without policing in equation (1) and the policing model in equation (2).

Firstly, in the non-policing model, z is designed to stay within the range from 0 to 1. It would be possible for an individual's z_{ig} to rise above 1, but this would be penalised as soon as z_g rose above 1, when everyone's fitness will be negative. In the policing model, however, z can increase to any level, and in fact the equilibrium level at $z = (1-r)/[r(1-k)]$ will often be much higher than 1 when r is very low. Frank tries to address this issue; how he does this is explained in section 3.10.

Secondly, the second part of equation (1), $1 - z_g$, is meant to represent the average group productivity, and in this equation it clearly does, because when you take the average of all the individual shares (which is the first part of the equation, z_{ig}/z_g), you get 1. In equation (2), the second square bracketed set of terms, $1 - (1 - a_g)z_g$, is what Frank calls "the local group's success". However, the individuals' fitnesses don't average out to be 1, they average out to be $1 - ka_g$, so the second set of terms is actually something other than the average group productivity. In fact it's what the average group productivity would have been had there been no cost to policing.

These three issues are not problems with how the model works, but taken together I think they make the model difficult to think about and difficult to justify. This is a motivation for developing alternatives in section 4.

3 Agent-based simulations of Frank's models

Frank doesn't feel the need to perform simulations of a population using his fitness equations. Running agent-based simulations provide a useful way to check the theory of policing models, and can be used to compare and evaluate competing models, especially when the gradients become complicated and difficult to visualise. This section describes a method for performing agent-based simulations which will then be used to explore Frank's policing models and some alternative ones.

3.1 General description of the simulations

A simple way of performing such a simulation is to create a population of agents, each with a value for the relevant traits being considered, and then repeat the following steps for some number of generations:

(1) Divide the individuals in the population into a number of fixed-size groups (assortment).

(2) Calculate their fitnesses using a fitness equation.

(3) Replace the population with a new generation of 'children' of the old population, where

  (a) the number of children of any individual is proportional to that individual's fitness, and

  (b) children have the same values for the various traits as their parents.

If this is repeated for a sufficient number of generations, and Taylor and Franks' reasoning is correct, individuals in the resulting population should have values for the traits that are close to the theoretical optima described in section 2.4.

3.2 Mutation

The algorithm described above doesn't include mutation. It assumes that no more variation will arise in the population, and the simulation must just choose a winning strategy from whatever it starts with. With a large enough population, and with each individual starting at a random value for their traits, this sort of simulation without mutation will probably still converge on values close to the theoretical optimal values — because something very close to the optimum will be available in the original population. It would be possible, however, for an ESS to be missed because there could be no starting individuals close to the optimum, or because all individuals close to the optimum could die out early on. Using mutation ensures that optimal individuals will eventually arise.

3.3 Assortment when individuals have only one trait

Assortment is the process of assigning individuals into groups. Simulating a population where individuals have only one trait that we are interested in, such as a population whose fitness is determined by equation (1), is straightforward. Firstly, individuals are ranked with respect to the trait, say z, and slotted into groups in that order. Perfect sorting ($r = 1$) occurs when individuals ranked 1 to n are in the first group, individuals ranked $n + 1$ up to $2n$ are in the second group, and so on. Following a suggestion from Marcus Frean, lower levels of r are achieved by first sorting all the individuals perfectly and then shuffling $1 - \sqrt{r}$ of them around randomly.

3.4 Assortment when individuals have more than one trait

Simply ranking individuals by a phenotypic trait does not work when there are two traits we are interested in, such as with a population whose fitness is determined by equation (2). For example, consider what would happen if we were to put individuals into groups based only on their z values. If we did this, then there is still a sense in which $\mathrm{d}z_g/\mathrm{d}z_{ig} = r$, because the mutation of an individual's z value will have an effect on which group he gets into. Suppose that z_{ig} changes to $z_{ig} + \Delta z_{ig}$. Then the chance that i will get into the 'correct' group for someone with that new z is equal to r, and the expected change in i's group, g, is going to be $r\Delta z_{ig}$. But such a system would be using a different notion of relatedness to Frank, who assumes that $\mathrm{d}z_g/\mathrm{d}z_{ig} = \mathrm{d}a_g/\mathrm{d}a_{ig}$, because it is clear that no change in a_{ig} will have any effect at all on a_g, no matter what r is.

I have developed a more 'genetic' approach to assortment, in order to follow Taylor and Frank's method, which is based on kin selection. Each individual has an extra tag, which represents its 'genotype', and which determines the values of its phenotypic traits such as a and z. Assortment is performed on the genotype. This provides a way of determining relatedness: two individuals are related if they have the same genotype. It might seem as if this all-or-nothing notion of relatedness is too simplistic, but in practice it is not a problem, because we are only interested in group relatedness. Group relatedness is easily calculated. The relatedness of group g is

$$r_g = \frac{1}{n_g^2} \sum_{t \,\mathrm{in}\, g} n_{tg}^2$$

where—

t stands for a genotype which occurs in group g; and

n_{tg} is the number of individuals in group g with genotype t; and

n_g is the size of the group.

So for example, if a group of ten individuals is made up of five of type A, and five of type B, the relatedness of the group will be $\frac{1}{10^2}(5^2 + 5^2) = 0.5$. Any group of identical individuals has $r = 1$, but a group of ten unrelated individuals has $r = 0.1$, and a group of 100 unrelated individuals has $r = 0.01$.

This genotype-based method has another advantage: relatedness is actually independent of the phenotypic values for each trait. This means that we have the possibility that lots of individuals in a group or population, for example, could have the same amount of policing or competitiveness as one another but still be genetically unrelated. This was not a possibility under the single trait ranking method of assortment described in section 3.3. ESS models assume that at equilibrium there will be a single phenotypic strategy that cannot be invaded; they say nothing about whether or not the population is genetically uniform, so these simulations are more flexible in this respect.

In addition, the original motivation for the policing trait was to solve the problem of how the negative effects of competition could have been suppressed in the first place, so a simulation that allows us to start off with a non-policing, competitive population is preferable.

3.5 An algorithm for genotype-based assortment

The assortment algorithm for multiple-trait genotype-based relatedness is similar to the single trait version in section 3.3: we first try to achieve near-perfect assortment (with r as close to 1 as possible), and then get a desired r value by randomly shuffling $1 - \sqrt{r}$ individuals around.

Achieving the original near-perfect assortment is more difficult when we are using genotypes, because we are no longer able to simply rank individuals with respect to a single value. Instead we want to achieve groupings that are as uniform as possible with respect to genotype. The algorithm starts off assuming that we have an array indexed by genotype, containing a count of the number of individuals in the population that have that genotype, or its frequency in the population.

(1) If any genotypes have frequency counts higher than the group size, allocate homogeneous groups made up entirely of individuals with that genotype.

(2) For each remaining genotype, in order from most frequent to least frequent, do the following while the frequency count for this genotype is greater than zero:

   (a) If a group has exactly the right amount of space in it to fit the remaining individuals of the current genotype, fill it up and move on to the next genotype.

   (b) If a group has more than enough space in it to fit the remaining individuals of the current genotype, allocate them to it and move on to the next genotype.

   (c) Pick the group with the most space in it and fill it up with individuals of the current genotype.

3.6 Relatedness, assortment and mutation

In general it is not possible to achieve exactly the desired level of relatedness in any simulation. Very low values for r are not always possible because every individual is always related to itself, so the mean r of a group of n different individuals will be $\frac{1}{n}$. The smaller the groups, the larger the minimum value of r will be.

There is also a problem with trying to achieve very high values for r, because it is often not possible to sort the individuals into groups perfectly. For example, suppose we have a population of twenty individuals with fifteen individuals of type A and five of type B. If we wanted to create four groups of five, then we could achieve $r = 1$ with perfect sorting. If we want to create five groups of four, however, there is no way to sort the individuals perfectly and we will have to settle for a mean r slightly less than one. This problem is partly just a consequence of the simplification that group sizes are fixed, but not completely; it will be a problem if we have any bounds at all on group size. There is no way, for example, to put unique individuals into a group with $r = 1$ unless we allowed groups to be of size 1, and a group of size 1 is not really a group.

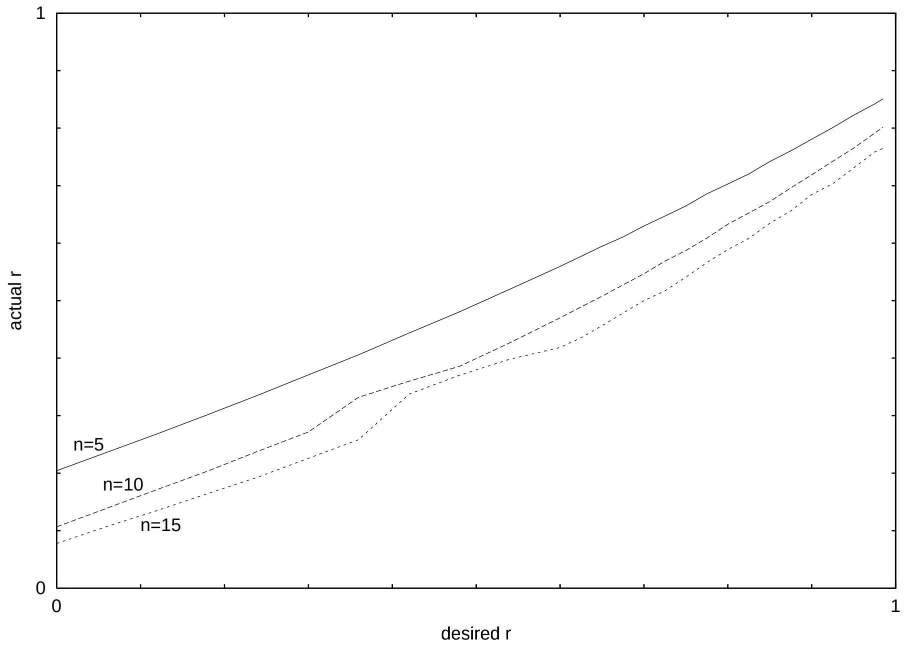

Figure 2: The relationship between the desired and actual values for r for three different group sizes(n). The maximum and minimum obtainable values of r vary with the group size.

It also turned out in the simulations that the maximum value of r that can be achieved is related to the number of mutations performed at each generation. Higher rates of mutation result in a more heterogeneous population which is harder to sort perfectly into groups. There is a tradeoff between the maximum level of relatedness achievable and the number of generations needed before the population converges to an equilibrium. Higher rates of mutation also mean that genotypes with optimal values for a and z will be available sooner, so the population gets to an ESS more quickly.

Figure 2 shows the relationship between the desired r and the actual measured value of r on typical simulations with three different group sizes. With small groups, very low relatedness is not possible, but high relatedness is easier to achieve than with larger groups.

3.7 The selection step

In the selection of the new generation at step (3) of section 3.1, the chance that each member of the new generation is of a particular genotype is proportional to the fitness of that genotype in the previous generation, as measured by the fitness equation. Fitnesses can go negative in Frank's model (see 3.10), so before the selection step, these are all made positive, by making the lowest fitness in the population positive, and shifting the other values upwards by the same amount.

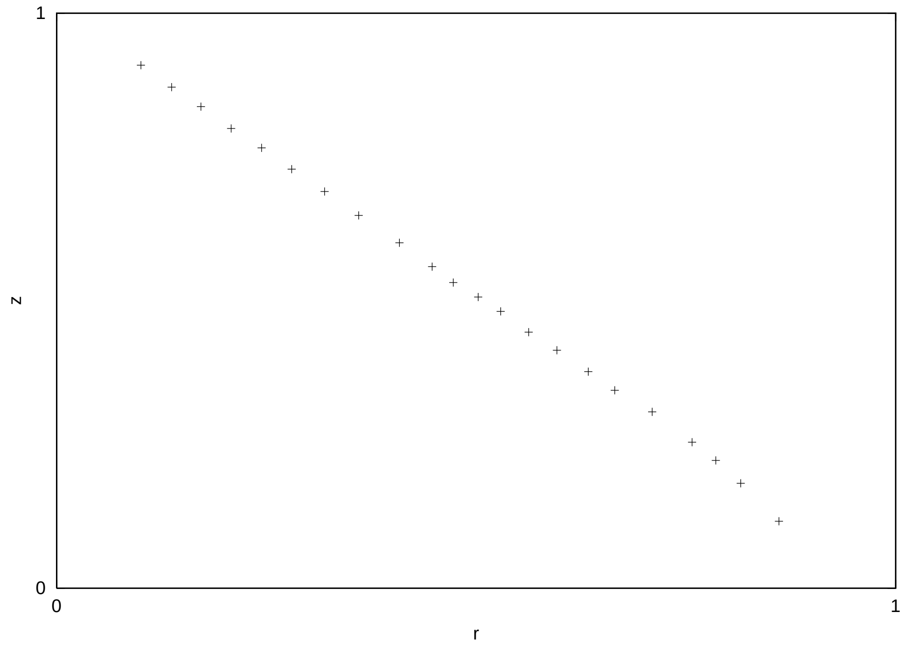

Figure 3: The results of simulations of Frank's non-policing model. Average competitiveness (z) is highest when group relatedness is low.

3.8 A simulation of Frank's simple model without policing

Figure 3 shows the results of a simulation of Frank's basic model without policing, described by equation (1). The general method described in section 3.1 was used, along with a mutation step between steps 1 and 2. At this step, a number of individuals are chosen at random from random groups. For each random individual, i, with genotype t_{old}, picked from group g, a new genotype t_{new} is created, and its z value is changed by a small amount from t_{old}'s z value[2]. Then all individuals in g (including i) that have t_{old} have their genotype changed to t_{new}. This type of simulation has the advantage that it makes explicit the fact that $\mathrm{d}z_g/\mathrm{d}z_{ig} = r$, because every time z_{ig} changes by say Δz_{ig}, z_g changes by (on average) an amount equal to $r\Delta z_{ig}$, because r other individuals in g share the same gene that i has.

The simulation was performed at a number of different desired values for r, and each one was left for 15,000 generations until the average value for z in the entire population stabilised. All individuals started out unrelated (they started out with different genotypes), but with $z = 1$. There were 105 groups of 11 individuals each. The y axis shows the mean value for z in the population, averaged over the last 500 generations of the simulation. The x axis shows the measured average r for all 105 groups in the population, again averaged over the last 500 generations.

The Taylor and Frank method for finding the optimal value, when applied to equation (1),

[2]Strictly speaking, step 2 of Taylor and Frank's method discussed in section 2.4 above is only reasonable if mutations have small effect (see [9]), so in these simulations the traits are only changed by a small random amount. The new genotype's z is picked from a Gaussian distribution around the old genotype's z, with a small variance. Because traits only change by a small amount at each mutation, longer simulations are required to reach equilibrium than would be needed if large mutations were allowed.

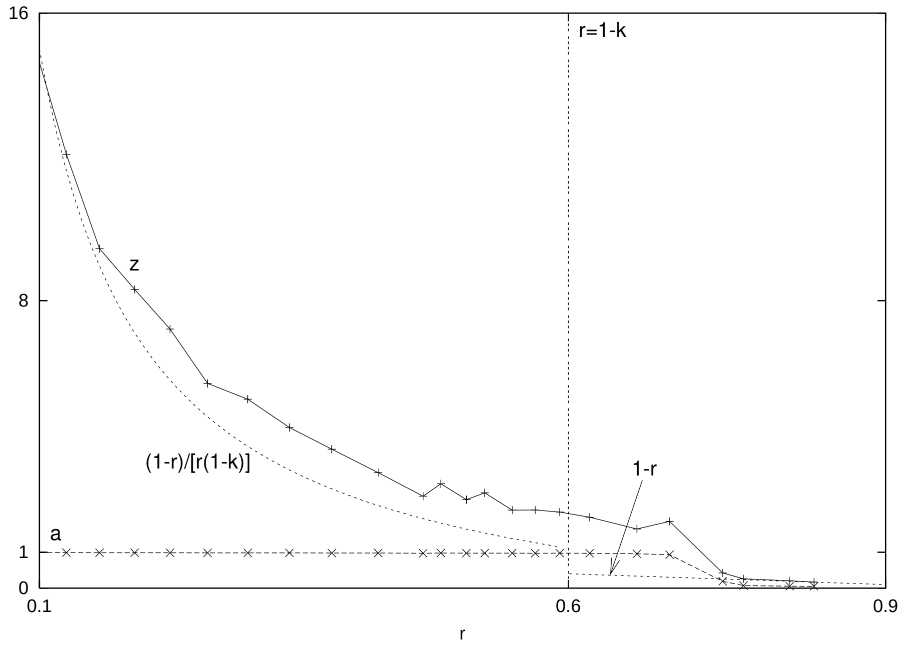

Figure 4: Average competitiveness and policing in a simulation of Frank's first policing model (equation (2) with $k = 0.4$. The theoretical optima for z are also shown.

---

predicts that $z = 1-r$. Figure 3 shows that the equilibrium points for z, from the simulation fall almost exactly along the $1-r$ line, in excellent agreement with the theoretical result.

3.9 A simulation of Frank's policing model

Frank's policing model of equation (2) was simulated using nearly the same method as in section 3.8. The main difference is in the mutations: a randomly chosen trait, either a or z, is mutated rather than always mutating z. There were 40 groups of 14 individuals each, and the cost of policing $k = 0.4$. In this simulation, and all the other simulations described below in section 4, every individual started out unrelated, with $z = 1$ and $a = 0$.

Figure 4 shows the results of a simulation of this policing model in equation (2). Also shown are Frank's theoretical optimal values for z, from section 2.5.

There are two main discrepancies between the theoretical values and the simulation values. Firstly, in the area where r is just over $1-k$, or between about $r = 0.6$ and $r = 0.7$ on the graph, things are overly competitive. The theory predicts that at equilibrium, a should be 0 and z should have dropped down to $1-r$; in fact, the simulation shows a up at 1, and z higher still. If the simulation gets into a situation where it has these values, we can see that it will be very hard to reach the equilibrium values, because the gradient $\partial w/\partial a$ is quite strongly positive when z is well above 1, but the gradient $\partial w/\partial z$ is only very weakly negative when a is close to 1. a is being pulled up more strongly than z is being pulled down, so it is difficult to reach the equilibrium point. Figure 5 shows this behaviour for a single run of a population when $r \approx 0.7$. At the start of the simulation, when both a and z are low, it only takes a small run of bad luck to move the population into the non-optimal state described

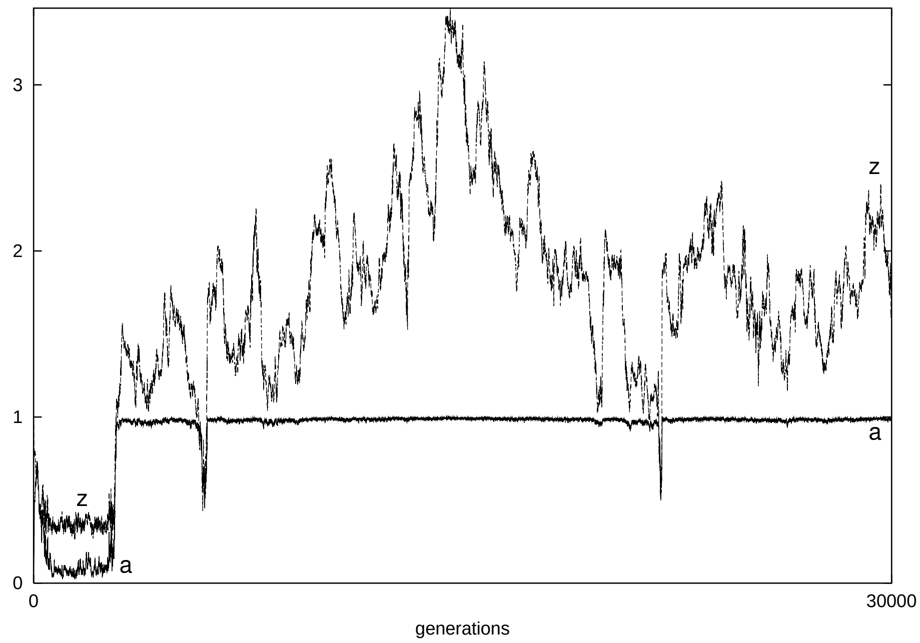

Figure 5: A single simulation of Frank's model when $r \approx 0.7$ and $k = 0.4$. Once z rises above a critical level, a jumps to 1, weakening the gradient for z, which wanders.

---

above, where z is higher than 1, but its downward pull is so weak that it is wandering; a's pull up toward 1 is stronger. The wandering behaviour of z is not helped by the fact that mutations in z can never go below 0, but can go above 1. The combined effects of mutations will therefore tend to pull the average z up slightly. This particular case illustrates that a simulation can show some of the features of a model that are not completely obvious from simply looking at the points where the gradients become zero.

The second discrepancy in the results in figure 4 is that the plotted simulation points for z are consistently higher than the optimal values when $r < 1 - k$. Although the gradient for z is very weak whenever a is up near 1, it doesn't look as if there's a good reason why z should remain consistently higher than the optimum, rather than just wandering to either side of it. I believe this is caused by the fact that the occasional mutant with a z that is very much higher than the optimum will do alright, whereas the occasional mutant with a z very much lower than the optimum will do very badly. So the high z mutants will last for a long time in the population, whereas the low z mutants will tend to disappear more quickly. This is especially true when the optimum for z is close to 1, but less true when the optimum for z is very high, at low r. The graph reflects this, because the z points plotted are the average for the entire population.

3.10 Frank's second policing model

A strange consequence of Frank's equation (2) model, which was mentioned in section 2.6 is that z must be allowed to grow above 1. Extremely high levels of competitiveness could make an individual's fitness go negative if they were introduced into a non-policing population, and Frank admits that "this type of incompatibility sometimes raises questions about a model" [3].

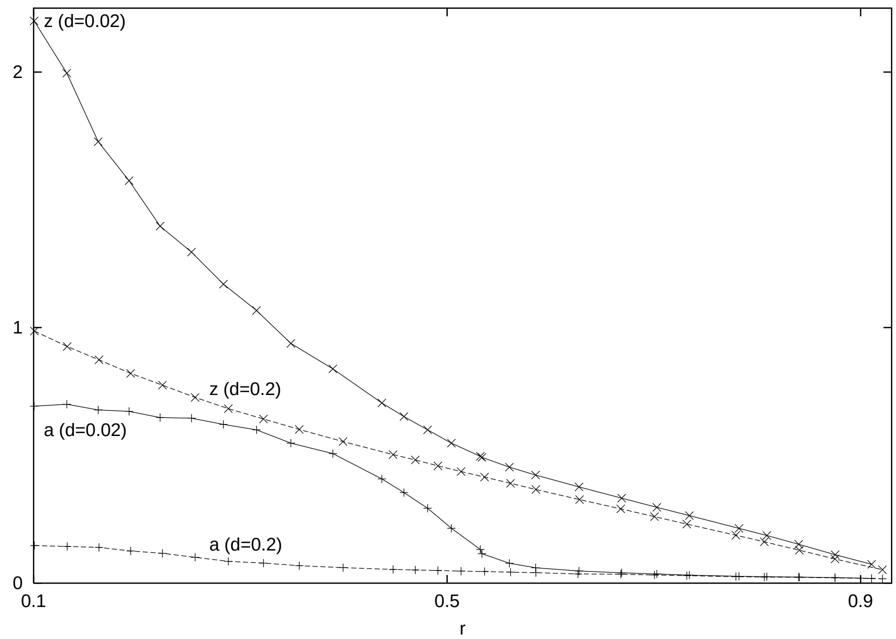

Figure 6: The results of simulations of Frank's second policing model for $k = 0.4$ and two different values for d. Lowering the cost of competition (d) also lowers the optimal level of policing (a), even when the cost of policing remains constant.

He claims that z increasing above 1 is not really a problem becuase it will only happen when there are extreme amounts of repressive policing, and that such high levels of competitiveness are just maladaptations to a non-policing environment. He also recognises that the model of equation (2) lacks realism because there is no cost to the competitiveness trait. In the real world, we would not expect to find that the average individual is five or six times as competitive in a society with policing than they would be in a society without policing, but these are the consequences of Frank's equation (2) model, where $z = (1 - r)/[r(1 - k)]$ at equilibrium.

Frank's solution to this problem is to add an explicit cost for competition, and offer us a new fitness equation even more complicated than the last one:

$$w_{ig} = (1 - ka_{ig} - dz_{ig})[a_g + (1 - a_g)\frac{z_{ig}}{z_g}][1 - (1 - a_g)z_g] \tag{6}$$

where d represents the cost of competition to the individual. Using this equation, it is possible (if you tweak d and k just right) to get z to remain below 1 at equlibrium even if it is allowed to climb higher in principle. This corresponds to the situation in which the addition of policing into a group doesn't affect the optimum level of competition.

Figure 6 shows the equilibrium points reached when equation (6) is simulated at two different values of d. In both simulations, $k = 0.4$. When $d = 0.02$, a settles at about 0.7 for low r, and z only climbs up over 2. When $d = 0.2$, a stays very low, because the cost of policing, k, is still the same, and z settles around 1. Clearly it would take quite a lot of fiddling around if you wanted both a and z to settle at around 1 for low r. This seems unsatisfactory.

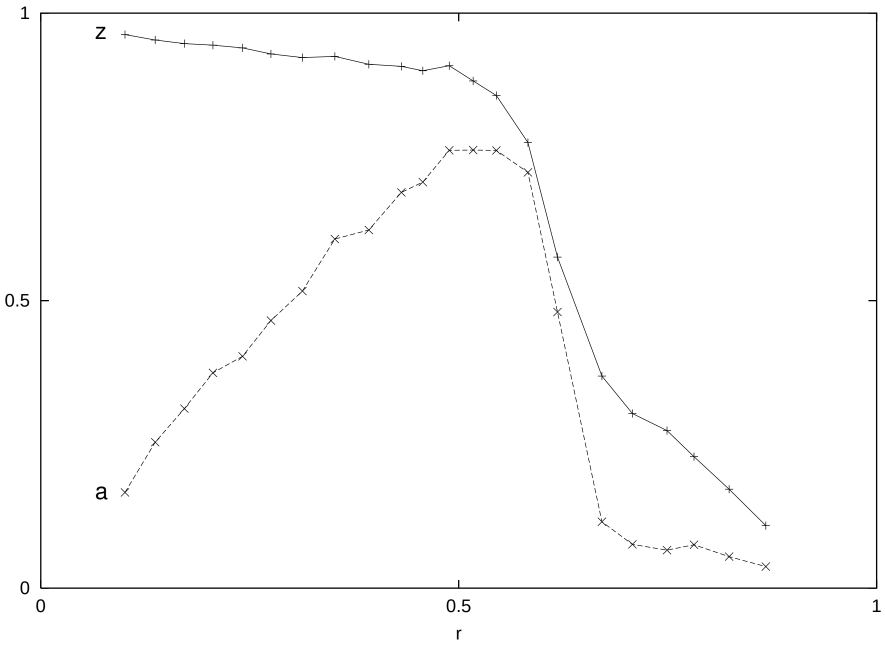

Figure 7: A simulation of Frank's first model of policing when the z trait is kept below 1, and $k = 0.4$. When z is not allowed to rise to its optimum level, very small decreases in competitiveness can result in an increase to policing.

One good thing about this equation is that a small change in r, when r is low, no longer means a huge changes to the z value at equilibrium, which is what we had with the equation (2) model.

3.11 Why z must grow above 1 in Frank's model

Figure 7 shows the results of a typical simulation of Frank's model of equation (2) when the z trait is kept under 1. This is done by not allowing any mutation that would take a genotype's z above 1. The a trait is being kept low at low levels of r, but increases as r increases up until the point when $r = 1 - k$. This is not really what we expect for any policing trait worthy of the name. Policing is supposed to be a trait that offsets the negative effects of group competition (or perhaps, in a model such as equation (6) with just the right values for d and k, suppresses competition). We expect it to be at its highest when group competition is at its highest. If policing is *lower* when the group is *more* competitive, we might not be able to solve the original problem of how evolution has produced so many cooperative groups, because presumably if anything like policing came about and thrived, it did so in a highly competitive environment.

3.12 The simulations are biologically implausible

The simulations are unrealistic from a biological point of view, because the mutation happens *after* the 'reproduction' and the assortment into groups. This rather strange ordering was used in the simulation to force the phenotypic derivatives, like $\mathrm{d}z_{ig}/\mathrm{d}z_g$, to be equal to r.

What it means is that the new mutants (who we can assume are all children of a single mutant in the previous generation) happen to all find themselves in the same group in the new generation, even when r is low. It is more realistic if we think of each generation of the program as actually representing several genrations of the individuals; it makes more sense that a number of mutants could arise in the same group over several generations.

4 Alternative policing models

Frank's policing models, even the modified one of equation (6) have a problem. It seems messy to have to use a model in which the competitiveness trait is allowed to rise to any level without restriction, but in which the policing trait is restricted to lie between 0 and 1. In this section I will outline some alternatives.

4.1 Linear interpolation using the endpoints of Frank's model

An alternative way of arriving at a policing model, suggested by Marcus Frean, is to extend Frank's non-policing equation (1) by trying to interpolate linearly between the two endpoints of the equation, that is, between what we would expect the fitness to be at the point where there is no policing ($a = 0$) and the point where there is full policing ($a = 1$). We assume that an individual's fitness when there is no policing, $w_{ig}^{a=0}$, will be exactly the same as it is in equation (1):

$$w_{ig}^{a=0} = \frac{z_{ig}}{z_g}(1 - z_g)$$

An individual's fitness under total policing, $w_{ig}^{a=1}$, will be the same as the fitness when everyone cooperates, minus a cost for the individual's investment in policing:

$$w_{ig}^{a=1} = 1 - k$$

A very simple fitness equation, if we just interpolate linearly between $w_{ig}^{a=0}$ and $w_{ig}^{a=1}$, is

$$\begin{aligned} w_{ig} &= a_g w_{ig}^{a=1} + (1 - a_g) w_{ig}^{a=0} \\ &= a_g(1 - k) + (1 - a_g)\tfrac{z_{ig}}{z_g}(1 - z_g) \end{aligned}$$

But here the group is paying the cost of policing rather than the individual. If we transfer the cost of policing to the individual, we can get an equation with the same endpoints above, but which is no longer simple linear interpolation:

$$w_{ig} = a_g + (1 - a_g)\frac{z_{ig}}{z_g}(1 - z_g) - k a_{ig} \tag{7}$$

$$\frac{\partial w}{\partial a} = rz - k \qquad \frac{\partial w}{\partial z} = 1 - r$$

The gradient for a shows that this model doesn't solve the problem that we had with 3.11. If we didn't like that model, we're definitely not going to like this one. It would be nice if a policing model had the property that the policing decreases as relatedness increases, becuase high relatedness fosters cooperation. The gradient above, however, doesn't have this

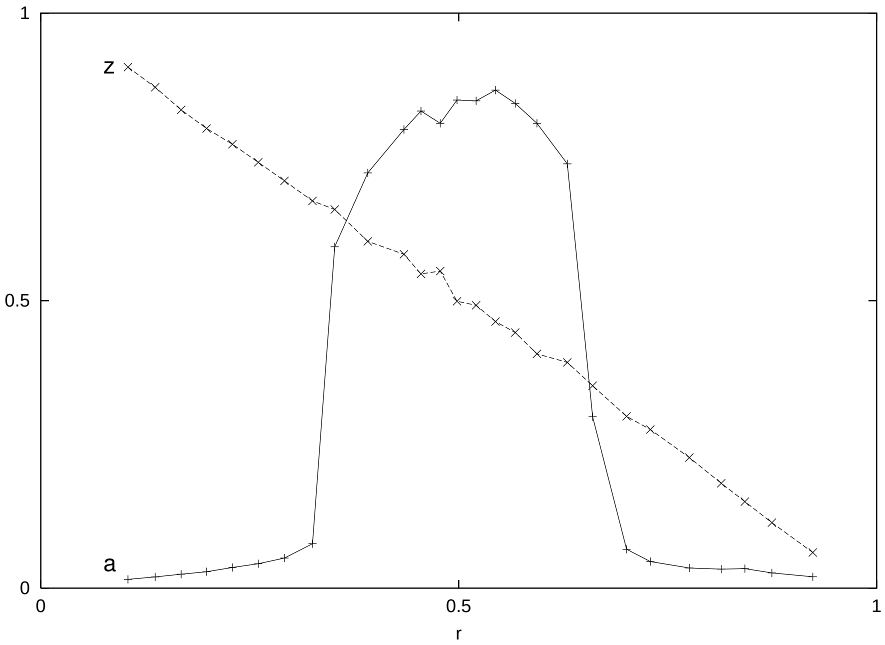

Figure 8: Simulation results for the 'linear' policing model when $k = 0.22$. Policing is at its highest for intermediate levels of competition.

---

behaviour. At equilibrium $z = 1-r$ and the gradient for a will effectively be $r(1-r)-k$, which will be lower for small values of r than it will be for medium values of r. All we have to offset this is the subtraction of k from the whole thing. This model has the unsatisfactory property that policing is less viable in more competitive groups than it is in averagely competitive groups. Figure 8 shows a graph of a simulation of the model of equation (7) for the case where $k = 0.22$, which confirms what we expect from the gradients for a and z. In fact it shows that the optimal level of policing is exactly the same for high r and low z as it is for low r and high z; the graph is symmetrical, so a is the same at equilibrium for both r and $1-r$.

4.2 The standard altruism model

The standard model of altruistic behaviour developed by evolutionary biologists uses a cooperation trait, x, rather than Frank's competitiveness trait, z [8]. The basic fitness equation is

$$w_{ig} = bx_g - cx_{ig} \tag{8}$$

where—

w_{ig} is the fitness of the i^{th} individual in the g^{th} group;

x_{ig} is i's level of cooperation;

x_g is the average level of cooperation in the g^{th} group;

	Player 2	
Player 1	C	D
C	$b-c$	$-c$
D	b	0

Table 2: Player 1's payoffs in a traditional prisoner's dilemma. C=cooperate, D=defect.

b is the benefit obtained by the group from cooperation; and

c is the cost paid by the individual for cooperation.

This equation corresponds to a simple prisoner's dilemma in which there is a cost, c, to be paid for cooperating, and a benefit, b, to be gained from the other player cooperating. The payoffs for cooperation and defection are summarised in table 2. The table shows the payoffs in a single move in a two-player game; the payoffs for an average round in a group game are summarised by equation (8). In the group game it is assumed that every player interacts with every other player in the group, round-robin style, so the individual's payoff depends on the average level of cooperation of all the other players in the group, x_g. The cooperative trait x can be thought of as the opposite of z in Frank's models, except that it really is restricted to the range between 0 and 1.

4.3 The addition of policing to the standard altruism model

A similar equation to (7) can be formulated using the two endpoints of the standard altruism model. An individual's fitness when there is no policing is the same as in equation (8), and fitness under total policing will be the same as the fitness when everyone cooperates, minus a cost for the individual's investment in policing:

$$w_{ig}^{a=0} = bx_g - cx_{ig} \qquad w_{ig}^{a=1} = b - c - k$$

where k represents the cost to the individual for policing. Again, we don't interpolate linearly between $w_{ig}^{a=0}$ and $w_{ig}^{a=1}$, but subtract the cost of policing from the individual rather than the group, and we have the same endpoints as we would with simple linear interpolation:

$$w_{ig} = a_g(b-c) + (1-a_g)(bx_g - cx_{ig}) - ka_{ig} \tag{9}$$

$$\frac{\partial w}{\partial a} = (b-c)r(1-x) - k$$

This gradient shows a similar problem to the one for equation (7). As r increases, x will also tend to increase, because cooperation is favoured when relatedness is high. The increase in r will be offset by a corresponding decrease in $1-x$, and the gradient will be at its highest for medium values of r. Again we have the strange property that there isn't as much pressure on policing at low r as there is at medium r.

	No of productive interactions	Payoff per productive interaction
Normal group	1	$bx_g - cx_g$
Policing group	$1 - a_g$	$a_g(b-c) + (1-a_g)(bx_g - cx_g)$

Table 3: The total group payoffs for policing and non-policing groups

4.4 The group should pay a cost for policing

Another problem with equations (7) and (9) is that the cost is only dependent on the individual's contribution to policing and not the entire group's. In equation (7), the problem becomes apparent when we consider the case $z_{ig} = z_g = 0$. The last term becomes 0 and the fitness is $w_{ig} = a_g - ka_{ig}$. Because there's no cost associated with a_g, a single individual with $a_{ig} = p$, say, gets $a_g - kp$, and is actually better off in a group that's doing a lot of policing than in a group that's doing only a little bit of policing, *even though everyone is cooperating.* There's a similar argument for equation (9), where at $x = 1$ the fitness is $w_{ig} = a_g(b-c) + (1-a_g)(b-c) - ka_{ig}$. It makes no difference whether the other members of your group are out there policing or not, because the group policing level is for free. There are good reasons to expect that this is very unrealistic: policing, in the sense that most people understand it, is only productive if there are defectors around to police. A fully cooperative group should suffer some penalty for wasting its energy on policing rather than using it more productively.

4.5 A model that treats policing as a lost opportunity to cooperate

If we think of policing and cooperating or competing as alternative actions available to any agent, then we get a cost for policing naturally. The cost for policing is just the lost opportunity to gain benefits through cooperation (or defection). In this model, time or resources spent on policing are time or resources than cannot be used for the productive interactions.

An interaction is an opportunity to either police, cooperate, or compete. For simplicity, we can think in terms of a prisoner's dilemma, in a single round of which every member of the group chooses either to police or to play the normal game against another random group member. The outcome of each interaction is affected by whether it is policed or unpoliced.

Policed interactions have two properties,

(1) They produce benefits as if cooperation had been enforced during the interaction, and

(2) The benefits from the interaction are evenly distributed throughout the group.

A mechanism to enforce property (2) is somewhat hard to envisage (we need to stop corrupt police from pocketing the group's rightful benefits), but this is in fact one of the assumptions used by Frank, and it could be something as simple as keeping an eye on your neighbours.

Using the simple model of section 4.2, an individual can expect to get $bx_g - cx_{ig}$ fitness points in an interaction which is not subject to policing. An interaction which is policed yields $b-c$ for each player, the same as if they had both cooperated, although these points go to the group, not to the players themselves. Over the whole group, the fraction of productive actions that are policed is a_g and the fraction that are unpoliced is $(1-a_g)$, so out of these productive interactions, the entire group's benefits will be

$$a_g(b-c) + (1-a_g)(bx_g - cx_g)$$

Because the policing itself is unproductive, there will be a_g lost interactions which gain 0 fitness points for the group; in other words, a group which does policing will have only $1-a_g$ of an interaction to play with for every interaction that a nonpolicing group would have. Table 3 summarises the difference between a group that does policing and one that doesn't. Clearly there is a cost involved to being a member of a group that does a lot of policing: there will be fewer group benefits available because there are fewer productive interactions.

Calculating the individual's payoff is harder. First consider the individual's payoff from all the policed interactions. Every policed interaction that occurs in the group delivers $b-c$ to be divided up evenly, so the individual gains a share from these interactions. The total number of productive interactions available to the group (policed and unpoliced) is $1-a_g$, and of these, a_g are actually policed. The individual's share from these is $(1-a_g)a_g(b-c)$.

Every unpoliced productive interaction delivers $bx_g - cx_{ig}$ *to the individuals involved in the interaction*, not to the whole group. The total number of interactions available to the individual (policed and unpoliced) is $1-a_{ig}$, and of these, only $1-a_g$ are unpoliced. The individual can expect a total of $(1-a_{ig})(1-a_g)(bx_g - cx_{ig})$ out of unpoliced interactions, and the fitness equation is just the sum of this and the share from the policed interactions above:

$$w_{ig} = (1-a_g)a_g(b-c) + (1-a_{ig})(1-a_g)(bx_g - cx_{ig}) \tag{10}$$

The gradients for a and x are:

$$\frac{\partial w}{\partial a} = (b-c)[(1-2a)r - x(1-a)(1+r)] \qquad \frac{\partial w}{\partial x} = (1-a)^2(rb-c)$$

The gradient for x must be positive when $r > c/b$, and must be negative when $r < c/b$. When $r > c/b$, x is driven up to 1. If this is the case, then $\partial w/\partial a = (b-c)(a(1-r)-1)$, which is always negative, so a will be driven down to zero. When $r < c/b$, x will be driven down to zero, and $\partial w/\partial a = (b-c)r(1-2a)$, which is zero when $a = \frac{1}{2}$. So we expect $x = 0$ and $a = \frac{1}{2}$ when $r < c/b$, and $x = 1$ and $a = 0$ when $r > c/b$.

Intuitively, $a = \frac{1}{2}$ makes a lot of sense when we consider the motivation behind the model; it can be thought of as the point at which half the individuals in the group are policing and the other half are trying to perform their productive interactions. Because there is no extra benefit for an interaction being policed twice rather than once, $a = \frac{1}{2}$ is the point at which all the productive interactions are being policed, but no more.

Simultion of equation (10) gives the graph in figure 9, where $c = 1$ and $b = 3$. Although the equation is formulated in terms of x, $1-x$ has been plotted on the graph so that the model

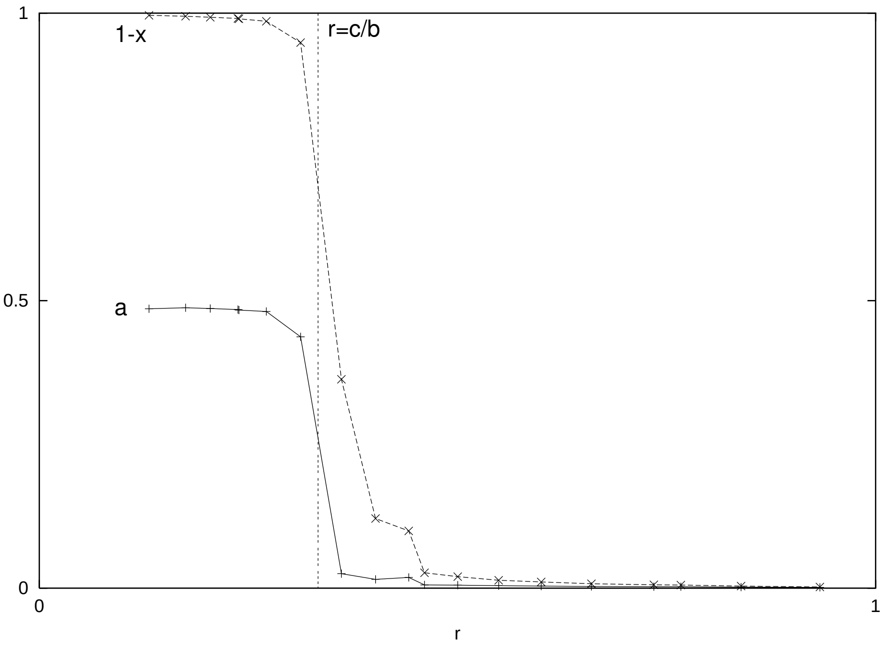

Figure 9: Simulation results for the 'lost opportunity' model of equation (10), with $b = 3$ and $c = 1$. Policing settles at $\frac{1}{2}$ when competition $(1 - x)$ is high.

can be easily compared to Frank's z. The results show the equilibrium value for a at around the theoretically optimal value $\frac{1}{2}$, for low r, and around zero for high r, as expected.

4.6 A similar model derived from an extended prisoner's dilemma

Following a suggestion of Marcus Frean's, I have tried to formulate the above model using pairwise prisoner's dilemma-type interactions and payoffs, in a game in which agents can choose to do a 'policing' action as well as the normal cooperation and defection of the standard prisoner's dilemma. Table 4 shows the payoffs for this extended prisoner's dilemma with policing. The entries for C (cooperate) and D (defect) in this table are exactly the same as for the traditional prisoner's dilemma, and the third column and row add policing. Whenever a policeman meets either a cooperator or a defector, both parties receive benefits of $(b-c)/2$. Effectively the policeman forces the other agent to cooperate, but is totally unproductive himself. The benefits from such interactions are shared equally between the two players (whereas in the above model the idea was that benefits from policed interactions go into a central pool). Because the policeman is unproductive, this is only half as good as if the policeman had cooperated, in which case both agents would have recieved $b - c$. Policing is also held in check by the fact that policemen are unproductive, and two policemen meeting one another both receive zero.

The expected fitness for an agent playing this game in a larger group can be written down straight from the table. An individual's expected payoff is the sum, for each entry in the table, of the payoff for that entry times the chance that the entry corresponds to the action taken, where the chance that i polices is a_{ig}, the chance that i cooperates is $(1 - a_{ig})x_{ig}$, and

	Player 2		
Player 1	C	D	P
C	$b-c$	$-c$	$(b-c)/2$
D	b	0	$(b-c)/2$
P	$(b-c)/2$	$(b-c)/2$	0

Table 4: The payoffs for player 1 when the prisoner's dilemma is extended to allow a third action: policing. C=cooperate, D=defect, P=police

---

the chance that i defects is $(1-a_{ig})(1-x_{ig})$. The equation can be written out in long form as follows:

$$
\begin{array}{rcccccc}
 & \text{Payoff} & \text{Player 1} & & \text{Player 2} & & \text{Table entry} \\
w_{ig} = & (b-c) & (1-a_{ig}) & x_{ig} & (1-a_g) & x_g & [\mathrm{C,C}] \\
- & c & (1-a_{ig}) & x_{ig} & (1-a_g) & (1-x_g) & [\mathrm{C,D}] \\
+ & (b-c)/2 & (1-a_{ig}) & x_{ig} & a_g & & [\mathrm{C,P}] \\
+ & b & (1-a_{ig}) & (1-x_{ig}) & (1-a_g) & x_g & [\mathrm{D,C}] \\
+ & 0 & (1-a_{ig}) & (1-x_{ig}) & (1-a_g) & (1-x_g) & [\mathrm{D,D}] \\
+ & (b-c)/2 & (1-a_{ig}) & (1-x_{ig}) & a_g & & [\mathrm{D,P}] \\
+ & (b-c)/2 & a_{ig} & & (1-a_g) & x_g & [\mathrm{P,C}] \\
+ & (b-c)/2 & a_{ig} & & (1-a_g) & (1-x_g) & [\mathrm{P,D}] \\
+ & 0 & a_{ig} & & a_g & & [\mathrm{P,P}]
\end{array}
$$

It turns out that this equation can be simplified down to

$$w_{ig} = (\tfrac{1}{2}a_{ig}(1-a_g) + \tfrac{1}{2}(1-a_{ig})a_g)(b-c) + (1-a_{ig})(1-a_g)(bx_g - cx_{ig}) \tag{11}$$

which is nearly the same as equation (10)! The difference between the two is in the first term, where equation (11) has $\frac{1}{2}a_{ig}(1-a_g) + \frac{1}{2}(1-a_{ig})a_g$ in place of $(1-a_g)a_g$. Equation (10) says, "every time someone in your group who is policing meets someone else who isn't, you get $b-c$"; your own gain from the shared group benefits isn't directly affected by your own personal level of policing. Equation (11) says, "if you are policing and you meet someone else who isn't, you get $(b-c)/2$, and if you aren't policing and you meet someone else who is, you also get $(b-c)/2$"; your share of the group benefits from policing is dependent on the difference between your level of policing and the average; the bigger the difference between you and the average, the bigger your gain from group policing benefits.

This affects how well an individual does when the population is very diverse with respect to a, but it makes no difference to the evolutionary stable phenotype. The gradients,

$$\frac{\partial w}{\partial a} = (b-c)(1+r)[\tfrac{1}{2} - a - x(1-a)] \qquad \frac{\partial w}{\partial x} = (1-a)^2(rb-c)$$

give the same results at equilibrium that we found for equation (10), even though the gradient for a is different. For $r < c/b$, $\partial w/\partial x$ must be negative, x is driven down to zero and $\partial w/\partial a = (b-c)(1+r)(\frac{1}{2}-a)$, which is 0 at $a = \frac{1}{2}$. For $r > c/b$, $\partial w/\partial x$ must be positive,

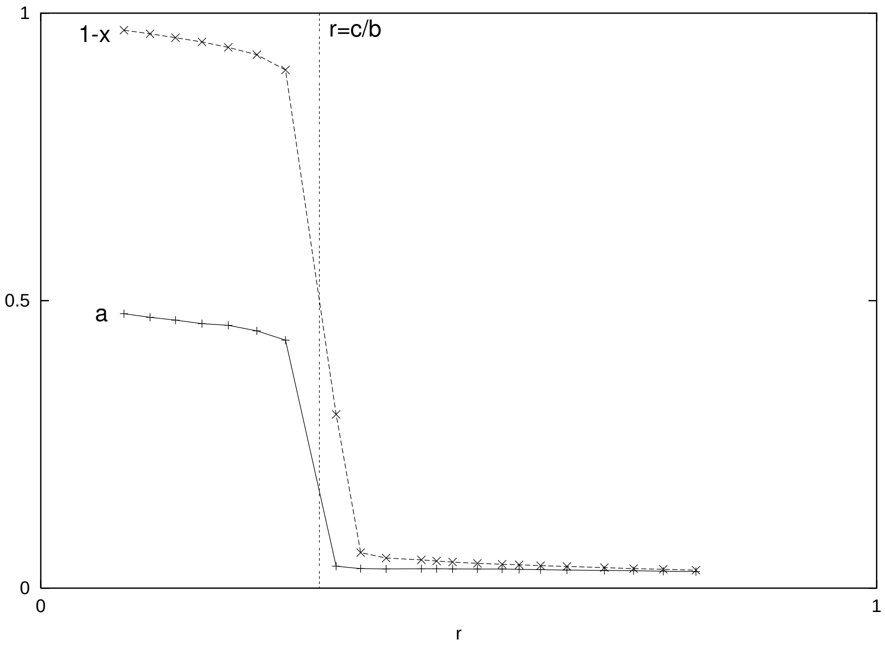

Figure 10: Simulation results for the extended prisoner's dilemma model of equation (11), for $b = 3$, $c = 1$.

---

and when $x = 1$, $\partial w/\partial a = -\frac{1}{2}(b - c)(1 + r)$, which is always negative, so a is driven down to zero.

The simulation of equation (11), again with $c = 1$ and $b = 3$, is shown in figure 10. There is a strong agreement with the theoretical results. As r approaches c/b from below, $1 - x$ drops slowly away from 1 and a drops slowly away with it from $\frac{1}{2}$, because the gradient for x is less negative as r rises towards c/b.

This model was formulated in terms of a two-player prisoner's dilemma, but the two-player prisoner's dilemma is in some ways an unsatisfying way to think about policing. Policing, at least in the way that it is presented by Frank, is usually thought of as something done by a group. In the prisoner's dilemma view it looks like policing is an action that a player can do to another player and automatically receive a reasonable benefit. The payoffs in table 4, however, are just meant to represent the average payoffs for the individual in a group that's playing the 3-action prisoner's dilemma.

5 Conclusions

There are a number of characteristics that could be used to evaluate policing models, and in the end all of them come down to what sort of concept of policing is plausible. In this section, I briefly discuss some of the important properties of the policing models explored in the project.

5.1 Does policing actually suppress competition?

Or does policing just remove competition's effect on fitness? The distinction here is partly semantic. Does competition refer to actual instances of competitive behaviour, or does it just refer to a desire or tendency to compete?

Of the models, it's only in Frank's second one, with the explicit cost of competition, that an increase in policing will bring about a reduction in the competitive trait z, although a 'reduction' only makes sense in terms of Frank's non-policing model, in which there is no explicit cost to competition. The cost of competition is totally intertwined with the cost of policing, and this makes the model complex. In order to stop competition going up (relative to the non-policing model) under policing, we have to set the two constants just right.

In Frank's first model of policing, policing does the opposite of suppressing competition — it actually increases it. It is difficult to think of a plausible situation in which policing would really have this effect, and the huge increase in competitiveness at low levels of r in this model almost rule it out in terms of realism.

In the 'lost opportunity' and extended prisoner's dilemma models, policing has no effect on the level of competition, it just moderates the effects of it. If we think competition actually refers to instances of competitive behaviour, then these models are also unrealistic.

5.2 Are levels of competition really unbounded?

In Frank's first model, there's no limit to how competitive you can be. This is an unpleasant property just from the point of view of neatness of the model, but I think it's also unrealistic to claim there's no limit on how competitive you can be while at the same time claiming that there's a hard limit ($z = 0$) on how cooperative you can be.

5.3 The prisoner's dilemma

The prisoner's dilemma type models provide a nice easy way to incorporate policing into a prisoner's dilemma game. This is an advantage. The prisoner's dilemma allows very low level simulations to be run with every action represented explicitly. However, it would be nice if there were a realistic explanation of what a policing action could realistically involve. It's a struggle to think of an example of behaviour that fits in with the payoffs of the two-person game in table 4. The extended prisoner's dilemma model is attractive because it expresses policing in terms of costs and benefits, and it is interesting to be able to compare it to cooperation and defection.

The 'lost opportunities' model doesn't represent a policing 'action' as such. Policing is abstract in the model, but the justification for the fitness equation assumes a world made up of pairwise interactions; it assumes that policing is divided up into 'units' that can affect exactly one other interaction.

5.4 Complexity

The alternative models from section 4 all have very simple, all-or-nothing behaviour. Policing is like a thermostat that turns itself on and off when competition reaches an intolerable level, and that level just depends on b and c. Frank's models are nicely complex — they provide a continuum of levels of policing. Of course, if a continuum is what we want, then we have to ask the question of whether all this complexity provides the best sort of continuum. It might be better to have, at equilibrium, policing increase exactly in tandem with competitiveness.

5.5 Do the models tell us how cooperative groups could have evolved?

The claimed motivation for dreaming up mathematical models of policing is that it will help evolutionary biology explain how cooperative groups could have evolved from competitive lower level units. It really depends on what level of explanation is required — these models are silent about the actual form that policing takes. Models only tell us that policing behaviour could be described in terms of a certain fitness equation with certain constants. Those constants are bound to be different for any real sort of policing behaviour, and it seems unlikely that they will be measurable in practice.

Acknowlegement

Marcus Frean, my supervisor, provided great advice on things to try, many hours of help with the maths, and helpful comments on two drafts of this report. Thanks Marcus.

References

[1] Richard Dawkins. *The Selfish Gene*. Oxford University Press, 1976.

[2] Steven A. Frank. Mutual policing and repression of competition in the evolution of cooperative groups. *Nature*, 377:520–522, 1995.

[3] Steven A. Frank. Policing and group cohesion when resources vary. *Animal Behaviour*, 52:1163–1169, 1996.

[4] William D. Hamilton. The genetical evolution of social behaviour. *Journal of Theoretical Biology*, 7:1–52, 1964.

[5] G. R. Price. Selection and Covariance. *Nature*, 277:520–521, 1970.

[6] F. L. Ratnieks. Reproductive harmony via mutual policing by workers in eusocial hymenoptera. *The American Naturalist*, 132:217–236, 1988.

[7] J. Maynard Smith and G. R. Price. The logic of animal conflict. *Nature*, 246:15–18, 1973.

[8] Elliot Sober and David Sloan Wilson. *Unto Others: the Evolution and Psychology of Unselfish Behaviour*. Harvard University Press, 1998.

[9] Peter D. Taylor and Steven A. Frank. How to Make a Kin Selection Model. *Journal of Theoretical Biology*, 180:27–37, 1996.

[10] David Sloan Wilson and Elliot Sober. Reintroducing group selection to the human behavioural sciences. *Behavioral and Brain Sciences*, 17:585–654, 1994.