Evolutionary Ecology Research, 2007, **9**: 1077–1096

A simulation study of the evolution of ageing

Owen G. Woodberry,* Kevin B. Korb and Ann E. Nicholson

Clayton School of Information Technology, Monash University, Clayton, VIC 3800, Australia

ABSTRACT

Hypothesis: Although ageing may reduce individual fitness, it could arise because ageing increases organism turnover and enhances genetic diversity within subpopulations.

The model: We have developed an ALife simulation environment of subpopulations of co-evolving hosts and diseases to test hypotheses, finding differing circumstances in which adaptive, and non-adaptive, ageing evolution occurs.

Conclusions: Our simulations demonstrate conditions that are favourable to the adaptive evolution of ageing. The adaptive evolution of ageing requires a reinterpretation of group selection theory, tying it strongly to what we already know about inclusive fitness theory.

Keywords: adaptation, ageing, co-evolution, diversity, group selection, inclusive fitness, kin selection, units of selection.

1. INTRODUCTION

Since Darwin proposed the theory of evolution by natural selection, observations of ageing in organisms have confounded biologists. The wide variety in longevity observed across species indicates that ageing is a species characteristic. Evolution by natural selection holds that many of the traits we observe in organisms are the result of adaptation to the environment of their ancestors. That is, selection between alternative alleles is based on reproductive success, or fitness. This is often, erroneously, interpreted simply as individual selection – the reproductive success of the individual organism bearing the gene. Clearly, following this interpretation, a self-destructive trait, such as ageing, can only have a negative effect on fitness. Organisms that age faster will have fewer expected offspring than others. The benefits of ageing, if there are any, are not benefits directly received by the individual. Although early biologists had postulated selection mechanisms that ascribed benefits to the ecosystem, species or group, such suggestions lacked theoretical backing until recently.

Evolutionary explanations of ageing were initially studied theoretically. Weismann (1889) was one of the first biologists to propose an evolutionary explanation for ageing. He postulated that ageing was an adaptive trait and ascribed its benefits to the species/group as removing older, and therefore worn-out and wasteful, organisms, thereby freeing resources for the young. A major criticism of Weismann's theory was its heavy reliance

* Author to whom all correspondence should be addressed. e-mail: owenw@infotech.monash.edu.au

on a hypothetical group selection. This led others to propose non-adaptive theories of ageing – those that proposed ageing was a by-product of another adaptive trait or a chance product. These alternative theories include Williams' (1957) antagonistic pleiotropy theory and Medawar's (1952) mutation accumulation theory. Recent experimental evidence creates a compelling argument that ageing is an adaptation (Mitteldorf, 2004), motivating a reinvestigation into the possible altruistic benefits.

A decade after these non-adaptive theories of ageing were proposed, the papers of Wynne-Edwards (1962) on group selection and Hamilton (1964) on inclusive fitness theory attempted to give mathematical backing to selection mechanisms beyond individual selection. Maynard Smith (1976) showed that the turnover of groups, via extinction and pioneering, renews the global population occasionally with purely altruistic groups. This ensures that cheaters, which may have the fitness advantage within groups, do not take over the population. Although the group selection concept has remained contentious, inclusive fitness theory, in the form of kin selection, has gained much credibility. We argue that group selection is supervenient on kin selection. That is, group selection is the result of isolated groups of organisms, which, over multiple generations, become internally highly related. Therefore, traits that benefit the group in competition with other groups will have a higher inclusive fitness, even if individually disadvantageous.

In this paper, after providing a review of the debate surrounding the evolution of ageing, we outline our own hypothesis of ageing for the sake of group diversity. We then discuss the mechanisms of group and kin selection, which allow the evolution of altruistic traits, focusing on the relationships between them. We describe the design of our ALife simulation environment with competing groups of co-evolving hosts and diseases. We then discuss our experiments with the simulation to test our hypothesis of the evolution of ageing in comparison with the popular non-adaptive explanations. We found that shorter life spans resulted in populations of hosts with greater genetic diversity that were more resilient to a disease population exploitative of genetic similarity. When the groups of hosts were pitted against each other, with mechanisms of group extinction and pioneering implemented, a positive selection for short life spans readily occurred. In simulations conducted with by-product genes enabled, we found that the by-product genes, which were selected for their positive effects on the individual, offset the benefits of the simple ageing genes.

2. THE EVOLUTION OF AGEING

Ageing is defined as the general deterioration of an organism, and its eventual death, by internal causes (Williams, 1957). The rates at which different species age is a perplexing phenomenon. On the one hand, a wide variety in longevity observed across species indicates that ageing is a species characteristic. On the other hand, the obvious fitness costs of ageing on individual survival would cause strong direct selection pressure against it, suggesting that ageing must have some other greater positive benefit or be a side-effect of a more essential characteristic (i.e. non-adaptive). Recently, compelling experimental evidence has accumulated that ageing is an adaptation (Skulachev, 1997; Bredesen, 2004; Mitteldorf, 2004).

Adaptive theories

Weismann (1889) was one of the first biologists to publish an evolutionary explanation of ageing, holding that ageing is an adaptation – that is, selected for its own sake. In

recognition of the obvious direct cost of ageing to the individual, Weismann argued that ageing death is beneficial to the species/group as it removes the worn out individuals that 'are not only valueless to the species, but they are even harmful, for they take the place of those which are sound' (Weismann, 1889). Weismann believed that over time an individual would be unable to avoid accumulating slight injuries, causing it to become defective and crippled, affecting its value to the species. This is contrary to the claims of other authors who hold that Weismann claimed the benefit of ageing was the adaptability of the species/group to a changing environment (Goldsmith, 2004; Mitteldorf, 2006).

Williams (1957), responding to Weismann's hypothesis, summarized the most forceful of the criticisms of the theory:

1. The fallacy of identifying senescence with mechanical wear.
2. The extreme rarity, in natural populations, of individuals that would be old enough to die of the postulated death mechanism.
3. The failure of gerontological research to uncover any death mechanism.
4. The difficulties involved in visualizing how such a feature could be produced by natural selection.

In the first criticism, Williams argues that Weismann's proposed benefit of ageing hinges on the assumption that the older members of the species will have deteriorated, regardless of senescence. He contends that Weismann's argument is circular – that is, it assumes what it claims to explain, which is declining vigour with age. The remaining criticisms outline potential difficulties with any adaptive theory of ageing, without, however, falsifying the thesis.

Overall, Weismann's hypothesis was heavily criticized, particularly for being unable to translate its proposed fitness benefits of ageing into individual organism fitness. Consequently, alternative theories have held that ageing has an inherent negative effect on fitness and therefore must be non-adaptive.

Non-adaptive theories

The theories of antagonistic pleiotropy and mutation accumulation are among the more popular theories to propose alternative evolutionary explanations of ageing that are non-adaptive. Antagonistic pleiotropy, proposed by Williams (1957), holds that ageing is a side-effect of genes that are selected for other beneficial properties. Williams assumes that (1) pleiotropic genes exist that have opposite effects on fitness at different stages of the bearer's life-cycle, and (2) that the probability of reproduction decreases with age, which follows from external causes of death and injury in the natural environment. He argues that pleiotropic genes that display benefits earlier in life could be selected regardless of self-destructive costs associated with the gene later in life. Mutation accumulation, proposed by Medawar (1952), holds that ageing is the result of mutational load. It is similar to the antagonistic pleiotropy theory in assuming decreasing selection pressure with age. Deleterious genes affecting the organism later in life will arise by mutation faster than they can be selectively removed from the population.

Mitteldorf (2004), Bredesen (2004), and Skulachev (1997) argue that, although these theories had good support 40 years ago, experimental evidence has since accumulated that is inconsistent with both. For example, pleiotropic theories, which require a direct causal link

between fertility and ageing, are undermined by experiments that increase longevity without any apparent cost to fertility (Leroi *et al.*, 1994). Also, caloric restriction experiments demonstrate that animals forestall ageing, without any fitness cost, in times of dietary stress (Weindruch and Walford, 1986), which raises the question of why ageing is not forestalled in times of plenty. This has led to a resurgence of research into the possible adaptive benefits of ageing, including Mitteldorf's (2006) demographic theory, Skulachev's (1997) phenoptosis theory, and our own theory, which we discuss next.

Diversity hypothesis

Our hypothesis, like that of Weismann, postulates an adaptive explanation of ageing – that is, it is selected for its own fitness benefit. We argue that ageing has a group fitness benefit that can outweigh the individual fitness costs. Groups with shorter individual life spans will turn over faster and consequently have greater diversity. In co-evolution scenarios, such as predator–prey and host–parasite interactions, groups with greater diversity will be less easily exploited, creating a stronger and healthier population. Similarly, Edmunds and Alstad (1978) proposed that an aphid colony, able to experience many generations on an individual tree, could evolve to thwart the defences of that tree.

In the following sections, we discuss mechanisms of selection that potentially favour altruistic traits and use them in the design of an ALife simulation of host–parasite dynamics. Guided by Williams' criticisms of the Weismann hypothesis (see above), we conduct experiments with our simulation, demonstrating the evolution of ageing and validity of our hypothesis in that particular simulation environment.

3. UNITS OF SELECTION

The evolution of an ageing trait would appear to be in direct conflict with the individual selection concept of natural selection. Individual selection refers to selection of the organism with the greatest individual fitness, measured as its ability to survive and reproduce, or simply the individual's expected number of descendants. It is easy to see that an organism exhibiting a trait, such as early ageing, which by definition reduces its own survivability, will leave fewer descendants than a competing organism without the trait, other things being equal. The benefits of an ageing trait, if there are any, could only be received by organisms other than the organism exhibiting the trait. For this reason, such a trait, if beneficial, must be altruistic. To give an explanation of such an altruistic adaptation, one must call upon a mechanism of selection that incorporates such altruistic benefits, or otherwise deny that it is an adaptation.

There have been two notable attempts to explain such a mechanism: group selection, attributed to Wynne-Edwards (1962), and inclusive fitness theory, proposed by Hamilton (1964). Group selection differs from individual selection in that it is the group rather than the individual organism that selection acts upon. Inclusive fitness differs from individual fitness by including the fitness effects accrued by relatives. This gene selection mechanism is often labelled 'kin selection'.

In the succeeding sections, we describe both of these mechanisms. We then consider the relationship between them, often thought to be a clear contrast. We argue that group selection is supervenient on kin selection, at least in most real-world scenarios.

Group selection

The group selection concept has an enduring bad reputation among many. The concept has been used in numerous arguments without due regard for theoretical backing. However, there have been working descriptions of group selection models, which, at least, demonstrate the logical possibility of the mechanism. Lewontin and Dunn (1960) were the first to give a workable description of a group selection model. In their model, it is the differing viability of the groups, together with their fecundity, that drives selection. Groups that contain selfish genes are more likely to become extinct and have less opportunity to produce emigrants to pioneer new groups. These types of models are reviewed by Maynard Smith (1976), who uses a simplified version of Gilpin's (1975) predator–prey model to describe them.

In Maynard Smith's model, a sharp distinction is drawn between selfish and altruistic groups, ignoring the transient groups based on the assumption that altruists will be quickly replaced by selfish free riders who share the altruist's benefit without paying any cost. Maynard Smith observes that the fate of the altruistic gene groups is dependent on the expected number of successful selfish emigrants from selfish groups during the lifetime of the group (Maynard Smith, 1976). A successful selfish emigrant is one that establishes itself and leaves descendants in a neighbouring empty or altruistic group. If the expected number of emigrants from selfish groups is greater than one, then the selfish groups will increase in frequency. Otherwise, the selfish groups will become extinct faster than they can found new groups and will therefore be selected out of the system.

Shortly after Wynne-Edwards' group selection mechanism was proposed, Hamilton (1964) introduced inclusive fitness theory. Because of the way in which these mechanisms were introduced, they are often seen as exclusive alternatives in explaining the establishment of altruism (Hamilton, 1975).

Kin selection

Inclusive fitness theory shifts the selection emphasis from the individual to the gene, be it held by the individual or, as a replica, by another – it is based upon the organism's individual fitness augmented by the harms and benefits caused to the fitness of neighbours, weighted by their relatedness (Hamilton, 1964). As the neighbours of the individual are also likely to be kin, the selection of the gene with the greatest inclusive fitness is often termed kin selection, a description coined by Maynard Smith.

Inclusive fitness theory was often cited as an alternative to group selection in explaining the evolution of altruistic behaviour until Price (1970) and Hamilton (1975) reformulated it into equations that identified different levels of selection: within- and between-group levels. Hamilton defines his groups unusually, as sharing *equally* the benefits of altruism, creating an important difference between his groups and real groups. As there is no preference for holders of the same genes, such as kin, altruistic genes are always selected against within the group, as any selfish free riders receive the same benefits as everybody else without paying the cost. This assumption is also made in the group selection models discussed above. For an altruistic gene to be positively selected, the magnitude of between-group selection must be greater than that of within-group selection. That is, groups with a higher frequency of altruists must perform better by increasing group fitness and hence increasing the relative size of altruistic groups in the population. If this increase outweighs the decrease in

frequency of altruists within each group, the gene will increase in global frequency. The more varied the frequency of altruists across the groups, and the more benefit bestowed by the altruists on the group, the greater this between-group selection will be.

Supervenience

Of the controversy over group selection and the importance of kin selection for the evolution of altruism, Hamilton (1975) remarked:

> Because of the way that it was first explained, the approach using inclusive fitness has often been identified with 'kin selection' and presented as an alternative to 'group selection' as a way of establishing altruistic social behavior by natural selection . . . . Kinship should be considered just one way of getting positive regression of genotype in the recipient, and that it is positive regression that is vitally necessary for altruism.

The idea is that inclusive fitness is more general than kin selection and might arise by mechanisms other than kin selection. What is needed for group selection of altruistic behaviour is the positive association between group fitness and the altruistic genes required by the Price equation. Such an association can arise by kin selection or by other means. But, however it arises, it will result in differential group fitness leading to a spreading of altruism.

In 2005, we conducted experiments on a group selection model of Mitteldorf (2006), isolating and removing the kin selection component by means of an adoption queue (Woodberry *et al.*, 2005). We found that the model was no longer capable of supporting the evolution of altruism, demonstrating the model's reliance on kin selection. This shows that group selection is supervenient upon kin selection in the simulations conducted – that is, there may be multiple possible ways of realizing group selection, but kin selection is one of those ways (Kim, 1993). Kin selection can lead to within-group selection for altruistic behaviour, so long as the groups are not Hamilton's groups, which miraculously share the benefits of altruistic behaviour identically across all group members [see Mascaro *et al.* (2001) for an example of simulated non-Hamiltonian groups]. And kin selection can also lead to between-group selection for altruistic behaviour, as our study demonstrated. But, as Hamilton avers, the association between longevity of groups and altruism could in principle happen otherwise, for example by chance or (at least in the case of artificial simulations) by intervention of a Designer. We believe that, although kin selection is not logically necessary for group selection, it is practically necessary: nature provides no alternative basis for differential group survival and reproduction. Hamilton's (1975) further suggestion that we reserve 'group selection' for situations where kin selection is inoperative would empty the term of all practical application.

4. SIMULATION DESIGN

To test hypotheses of ageing, we designed a multi-agent ALife simulation environment. Agents interact within groups sharing a food source and reproducing, potentially sexually. Groups may also interact with each other by mechanisms of group extinction, via disease epidemics, and migration (pioneering new groups), catering for a group selection mechanism (see Section 3). The simulations used to test our diversity hypothesis have groups that contain co-evolving populations of host and disease agents. Host agents interact with each other, through mating, and with the disease population, which parasitizes them. The success

of disease transmission between hosts is determined by the genetic similarity of the host chromosomes. When designing simulations, it is necessary to consider the trade-off between the complexity of the model and its completeness. Although more complex models are harder to analyse, simpler models could neglect important mechanisms that allow validation against real systems (Grimm *et al.*, 2005). In our design process, we tried to find an optimal resolution for our simulation.

World and time

The simulation consists of a 10×10 grid of group locations, with the edges connected and forming a torus-shaped world. Each of the group locations may or may not contain a group of agents. At this level, there is a group selection mechanism with group extinction and migration. The advents of group extinction and migration are controlled by within-group factors – primarily the group population size. When an extinction occurs, the group contents are reset and it becomes available for a pioneering population. When a group produces a migration party, two agents are randomly selected from the population and a destination group location is randomly selected from the immediate 3×3 group block neighbourhood. In the interests of simplicity, agents are not permitted to migrate to already occupied group locations, 'invading' existing groups. Instead, the destination group is selected from the unoccupied neighbouring groups. Time is represented by cycles. During each cycle the world is updated, with statistics collected every epoch, usually 10 or 100 cycles.

Groups

Each group contains a food store, and a host and disease agent population. Energy within groups, as in real-world ecosystems, is closed – energy can enter the group via food production (e.g. sunlight) and only exits via agent living overheads (e.g. heat). When an agent within the group dies, any energy it holds is recycled *indirectly* back into the group food store, via food fertilization. This is important so as not to penalize groups that have a higher agent turnover. Such groups would otherwise lose the energy held by the dying agents, which would have an undesirable influence on the population size. The host population feeds directly from the food store and the disease population, if one exists, feeds on the host population. Each cycle the food store, recycled from the previous cycle, is incremented by an amount determined by the normal distribution $N(100, 10)$. The food is then divided equally among the host population, before all the agents – first diseases and then hosts – are updated in a random order.

Host population

The host population is the focus of the simulation, and it is this population that determines when extinction and migration occur. Group extinction occurs when the host population completely dies out, which is usually caused by a disease epidemic. At the beginning of each cycle every group is tested, in random order, for the production of a migration party, determined by a probability proportional to the size of the host population, usually, $pr = population\ size \times 10^{-3}$. Host populations that have greater resistance to disease transmission and infection will be larger in size, since less energy is diverted into sustaining the disease population.

Agents compete in a single location, and thus on a purely individual selection basis. A kin selection effect could be catered for by providing a spatial structure within the groups, where kin are more likely to affect each other, but this was considered unessential at this level, and was omitted for the sake of simplicity.

Host agent

Each host carries a chromosome inherited at birth, which is used to determine the genetic characteristics of the agents, covered in the following discussion. The age, health, and fertility of the hosts are updated each cycle and are used, in conjunction with relevant genetic characteristics, to determine the behaviour of the carrier, in any given cycle. The host agents have four different actions, some or all of which can be carried out during any cycle: migration, eating, mating, and dying. The migration action is tested for every group at the end of the cycle, after all agents have been updated, as per the discussion above.

Host eating

Every cycle, each host is allocated an equal share of the group's food store. This food is used to replenish the host's health and fertility stores. Health represents the agent's life force – that is, while positive the agent is still alive, and fertility is used for parental investments and to test maturity. The host's health, initialized at 10 units at birth, is decreased by a 1 unit overhead every cycle. Usually, food energy is diverted directly into the health store; however, in antagonistic pleiotropy experiments, energy division between health and fertility stores is instead determined by a gene. When either of the stores exceeds its capacity of 10 units, excess food is first diverted to the other store and then, if any food remains, it is recycled back to the group food store. After all the host agents have eaten, the disease population, if there is one, is updated. The disease population implementation is discussed later.

Host reproduction

Next, the host agent attempts a mate action. The hosts reproduce sexually but are genderless. First, the host agent is checked for maturity – that is, if its fertility store exceeds the amount required for parental investment (> 5 units). If the host agent is mature, it randomly selects a mate partner from its group, which is also tested for maturity. If both agents are mature, they then conceive an offspring, with an inherited chromosome and health initialized at 10 units, which is deducted, evenly, from the parent's fertility stores. The child is then added to the group population and acts independently in the next cycle.

Host death

At the end of the cycle, the host agent is tested for death conditions. There are four different causes of death. The agent dies:

- if an accidental death occurs;
- when it exceeds its genetic expiry age;
- in the event of an age-triggered death mutation; or
- if it finishes the cycle with zero or negative value in its health store.

Accidental death can occur at any time during the agent's life span, at a fixed probability of 0.1, each cycle. The genetic expiry age and death mutation causes of death are both triggered by age. A distinction is made as the former is caused by a simple ageing gene mechanism and the latter by a mutation accumulation mechanism (both of these mechanisms are discussed in greater depth later). For convenience of implementation, there is a maximum age of 200, which, due to the high rate of accidental death, has a negligible effect on death rate. Health deaths can occur due to overcrowding, disease epidemics or failure to maintain the health store (which only occurs when antagonistic pleiotropy genes are enabled). When the agent dies, it is removed from the population and its health and fertility are recycled back into the group food store for the next cycle, for reasons discussed above.

Host chromosome

The host chromosome contains components of: a genetic expiry age gene; an antagonistic pleiotropy gene; a mutation accumulation gene string; and a vulnerability gene string. Components may be evolving, fixed or inactive depending on the experiment conducted.

Genetic expiry age component

In all experiments the genetic expiry age gene, which is the focus of the study, is always active, either fixed or evolving. This gene is used at conception to determine an expiry age for the agent by sampling a normal distribution with variance proportional to its magnitude, i.e. N(*expiry age, expiry age*/3). The gene is inherited from one of the parents, determined randomly, and mutated by a normal distribution with variance proportional to its magnitude, i.e. N(*expiry age, expiry age*/100). As this gene has no side-effects, it is expected that fast ageing would always be selected against unless the scenario provides ageing its own selective value.

Antagonistic pleiotropy component

When antagonistic pleiotropy is enabled, a gene is used to determine the ratio of food energy directed towards health (maintenance) and fertility. This gene is inherited from one of the parents, determined randomly, and mutated by a normal distribution, i.e. N(*pleiotropic ratio*, 0.01). In most simulations, the rate of health death is negligible, due to the high rate of accident deaths. Thus any health death is considered an ageing death, as the agent has failed to maintain itself. However, we expect a synergy effect with disease epidemics as a population with the pleiotropic genes will be, on the whole, less healthy and therefore more prone to epidemic extinction.

Mutation accumulation component

When mutation accumulation is enabled, chromosomes contain a string of death genes (with values of either 'on' or 'off'). Every cycle, the death gene, at the index given by the agent's age is considered, and, if 'on', death occurs with a 0.5 probability. The string length is 200 bits, a gene for each age up to the maximum age. The gene string is inherited by crossover of the strings of both parents, with a 0.001 probability of mutation of each locus

along the string. This component has a similar, age-triggered death, effect as the genetic expiry age gene. However, unlike the simple expiry age gene, the positive selection of this gene is dependent on the probability of reproduction decreasing with age (see Section 2), which occurs in our simulations as there is a high rate of accidental deaths.

Vulnerability component

For diversity experiments, the agent chromosome contains a vulnerability string, used to test for disease infection, the details of which are discussed below.

Disease agent

The disease population is only enabled during the diversity simulations. The methods of disease infection and host vulnerability were adapted from those of Epstein and Axtell (1996). The disease agents parasitize the host population using host health for the production of disease offspring. The disease agents consist only of a chromosome containing an infection bit string, which is used to determine whether the disease successfully infects a host with a vulnerability bit string (discussed below). The disease, which only lives for a cycle, does not store energy; instead, energy is used immediately to cover overheads and reproduction.

Disease eating, reproduction, and death

The disease agents are updated each cycle producing, asexually, offspring that are transmitted to neighbouring hosts within its group. The number of disease offspring is determined by a split rate, with a default probability of 0.15, generating two if a split occurs, otherwise one. The disease offspring immediately absorbs its energy requirements from the new host's health store, decreasing it by 1 unit. If the host already has a sub-zero health, the disease transmission fails. Further infection is then determined (discussed in detail below) and, if successful, the offspring disease is saved for the next cycle. At the end of the cycle, the parent disease is resisted by its host and dies.

As disease agents are not transmitted between groups, a new disease is generated for each group at the beginning of every cycle to ensure that the disease population is never completely eradicated.

Disease infection

The success of infection is determined by interactions between the host vulnerability (lock) and disease infection (key) bit strings (see Fig. 1). The host vulnerability is 50 bits in length. It is held in its chromosome and is determined by uniform crossover of its parents' vulnerability strings, with, usually, a 0.01 probability of mutation flipping each bit value. The disease infection bit string is 5 bits in length. It is held in the disease chromosome and is determined by copying the parent infection bit string, with a 0.01 probability of mutation flipping each bit value. The success of infection is determined by testing whether the infection string matches a sub-string of the vulnerability string, aligning the ends. Hence, the more genetically similar host agents are within a group, the greater the chance of successful transmission. The host vulnerability mutator determines the diversification rate – that is, the diversity generated by each population turnover. Ideally, we would like to have the mutator

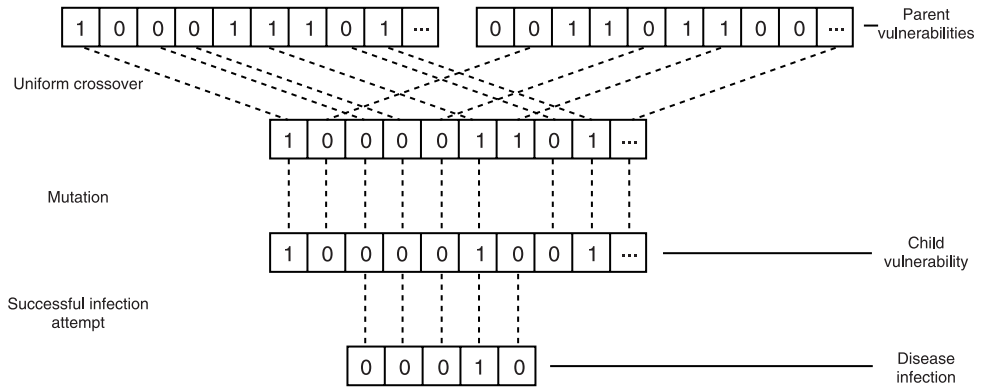

Fig. 1. The child's vulnerability bit string is determined via uniform crossover of the parents' vulnerability strings, with a chance of mutation flipping each bit value. An infection attempt is successful when the disease's infection bit string matches a substring of the host's vulnerability string.

itself evolve, allowing the rate of diversification to evolve. However, this was not possible due to the simplistic nature of the simulation, as there is no mechanism to account for the costs of random genetic leaps in design space, as there would be in the real world.

5. EXPERIMENTS

In this section, we describe experiments conducted with our simulation to test our diversity hypothesis of ageing. First, we examined the within-group dynamics of the host and disease populations. This allowed us to identify the relationship between ageing and diversity and thus the ability of the host populations with fixed ageing to resist exploitation by the disease population. Next, we conducted simulations with different combinations of the components of group selection, diseases, antagonistic pleiotropy, and mutation accumulation enabled. This demonstrates a variety of possible worlds with these different, but not mutually exclusive, hypotheses of ageing.

Experiment 1

Initially, experiments were conducted to measure the effects of varied genetic expiry age (5, 20, 50) and vulnerability mutation rate (0.000, 0.005, 0.010, 0.015) on group relatedness. This was done by conducting simulation runs with the disease component disabled, measuring group relatedness at the end of 5000 cycles (see Fig. 2). A relatedness of 0.5 means a perfectly mixed group, whereas 1.0 is interpreted as genetic homogeneity. Twenty simulations were run per parameter set. The group relatedness, the inverse of diversity, was measured as the average number of vulnerability string locations that are the same (i.e. inverse Hamming distance) between two randomly selected hosts within a group. Simulations were not conducted with very low genetic expiry ages, as host agents required at least a short amount of time to reproduce themselves. Figure 2 shows that groups with lower genetic expiry ages (i.e. greater population turnover) and a greater vulnerability mutation rate had, on average, lower group relatedness.

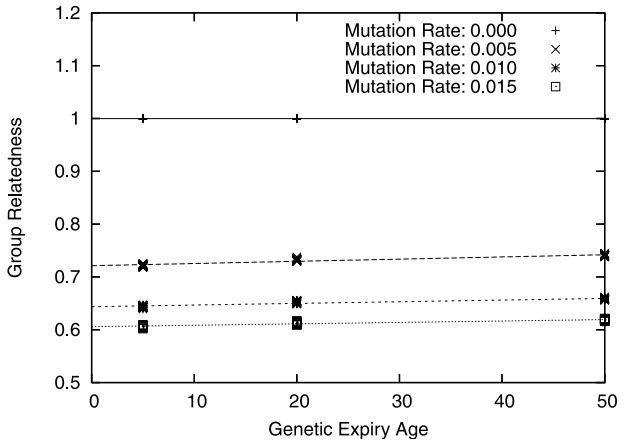

Fig. 2. The group relatedness resulting from varied genetic expiry age and vulnerability mutation rate.

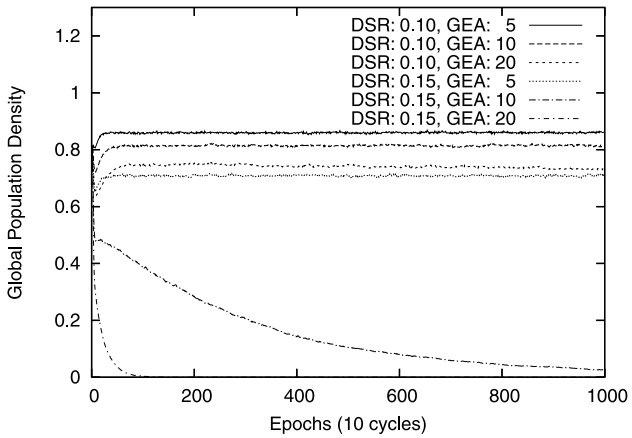

Fig. 3. The global population density plotted over time, for varied disease split rates (DSR) and genetic expiry ages (GEA).

Experiment 2

Next, simulations were conducted to gauge the effects on the host populations of varying disease split rate (0.10, 0.15) and genetic expiry age (5, 20, 50) (see Fig. 3). Global population density is affected by two factors: group population sizes and occupancy of group locations. Twenty simulations were run per parameter set. These simulations were conducted with the group migration component disabled, allowing the groups to become, and stay, extinct when disease epidemics caused them to. In these simulation experiments, and all remaining experiments, the vulnerability mutation rate was fixed at 0.01, meaning that the group diversity was determined by genetic expiry age. It was expected that groups with greater diversity (i.e. lower genetic expiry age) and lower disease split rates would be less susceptible to disease infection, and hence disease epidemics. In simulation runs conducted with greater

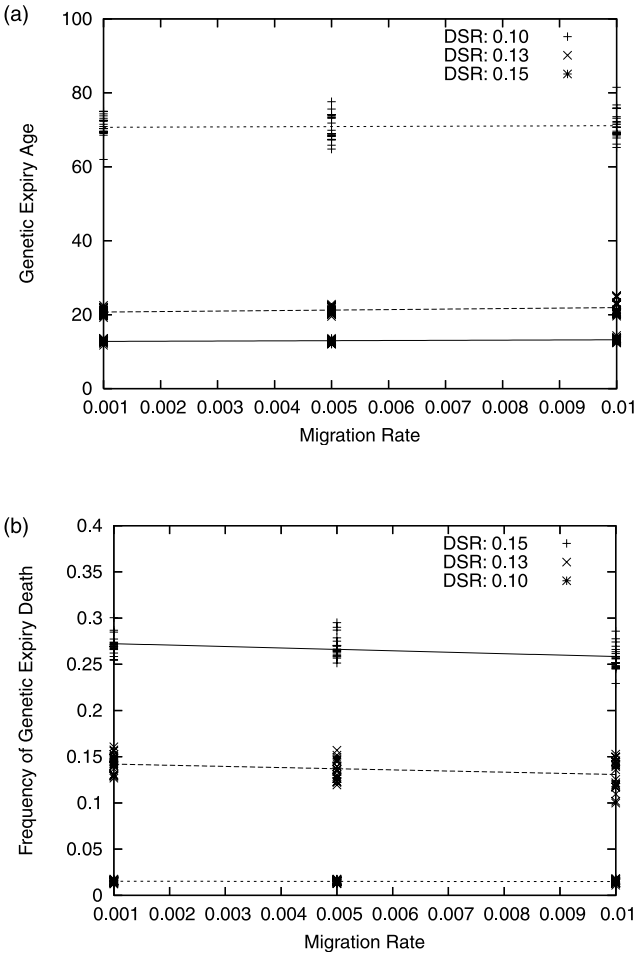

Fig. 4. (a) The evolved genetic expiry age and (b) the corresponding proportion of all deaths attributed to genetic expiry plotted against migration rate and disease split rate (DSR).

disease split rates and/or genetic expiry ages, groups become extinct and the global population is eventually wiped out. In simulation runs with lower disease split rates and/or genetic expiry ages, groups resist extinction, and establish equilibrium with the disease population with varying levels of infection.

Experiment 3

Finally, the migration rates were enabled, facilitating the expression of a group benefit of ageing, and the genetic expiry age gene was permitted to evolve. Simulations, 20 per parameter set, were run for 50,000 cycles and the resultant average genetic expiry age at the end of each run was recorded. Figure 4(a) plots the final genetic expiry age against varied migration rate (0.001, 0.005, 0.01) and disease split rate (0.15, 0.13, 0.10). Figure 4(b) plots the corresponding proportion of all deaths attributed to the genetic expiry age gene. The

migration rate affects how often the groups attempt to reproduce themselves and the disease split rate affects the chances of a disease epidemic driving the group extinct. It was expected that lower migration rates and greater disease split rates would be more favourable to the positive selection of the ageing gene. In simulations with greater migration rates and lower disease spilt rates, a greater genetic expiry age evolves. In some cases, the proportion of deaths attributed to genetic expiry, as can be seen from Fig. 4(b), becomes so low that there is only slight selection pressure for longer life spans. If migration rate were to be reduced further, group selection would fail, as groups would be unable to replace themselves before extinction. Similarly, if disease split rate were to be increased further, groups would suffer disease epidemics causing extinction before being able to replace themselves. For all the following experiments, the disease split rate was fixed at a probability of 0.15 and the migration rate at 0.001.

Calibration discussion

We observed a great sensitivity to variation in the vulnerability mutation rate and disease split rate, whereas the results were relatively insensitive to variation in migration rate (Fig. 4a) and genetic expiry age (Fig. 2). The insensitivity of the migration rate is a consequence of the implementation of group migration, where migrants may only relocate to empty group locations. Consequently, when all neighbouring group locations are occupied, the migration attempt fails. It would be expected that if the simulation were more complex, with migrants permitted to invade existing groups, it would be more sensitive to the rates of migration. In Fig. 2 group relatedness was plotted for varied fixed genetic expiry ages and vulnerability mutation rates. It can be seen that the vulnerability mutation rate need only be altered slightly to get the corresponding difference in relatedness of varied genetic expiry age. If the vulnerability mutation rate is decreased or increased too much, genetic expiry age will cease to evolve. This is because genetic similarity will be either too great, allowing diseases to quickly overrun host agents, or too small, with diseases unable to gain any foothold. A similar problem occurs with disease split rate; the difference in relatedness for varied genetic expiry age is quite small (see Fig. 2). Because of this, small changes in disease production also greatly affect the simulation with either an unsustainable host or disease population. There is a delicate balance between the two different co-evolving populations. These sensitivity problems are caused by the simplistic nature of the simulation and might be addressed by making the simulation more complex. In this work, to simplify the difficult simulation of co-evolving populations, disease restraint was not properly modelled and could not evolve, whereas in real-world situations, an equilibrium might be expected to evolve. Disease agents that overrun the host population (its energy source) would quickly drive themselves extinct, and would evolve restraint, via group selection.

Figure 5 maps the benefit of ageing and the relationship of some other simulation parameters. The ageing affects the group by increasing the rate of agent turnover, which, coupled with vulnerability mutation rate, generates group mutation and thus increases group diversity (decreases group relatedness). Groups with increased diversity will better resist the invasion of diseases, which are generated according to the disease split rate, and will be greater in size. The fitness benefit of ageing is expressed in the group size, which, together with the migration rate, determines the rate at which new groups are pioneered and whether the group will become extinct. Note, group size will also feed back into the production of mutations, as larger groups will generate more offspring every cycle.

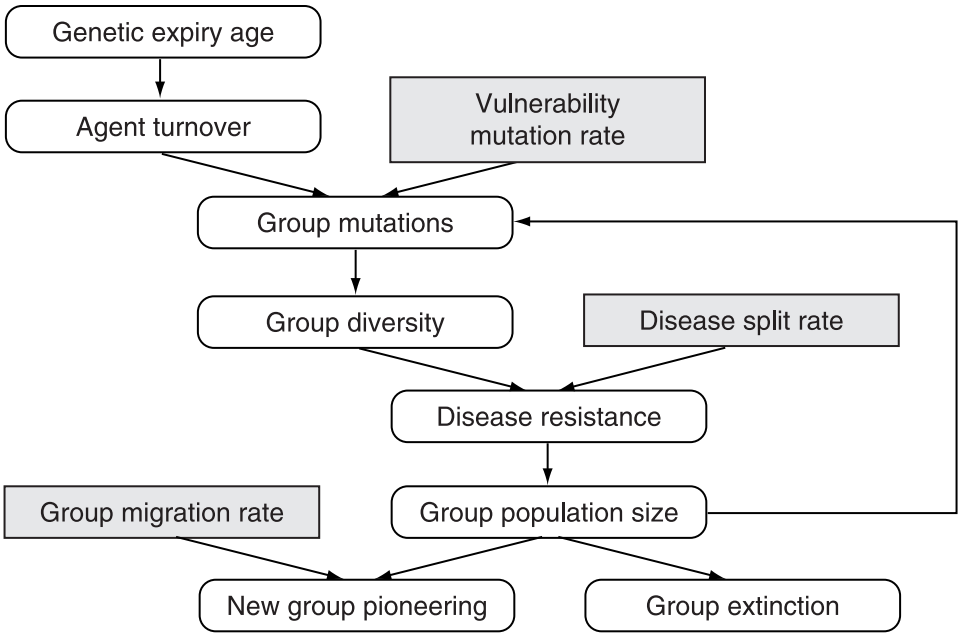

Fig. 5. Mapping of the causal pathway between genetic expiry age and group benefits of new group pioneering and delayed extinction (simulation calibration parameters shaded). Edges only indicate effect and not the type of effect. Genetic expiry age is fixed in calibration experiments to determine effects on variables down the causal pathway.

Experiment 4

We tested three different ageing hypotheses:

- our diversity hypothesis (DIV);
- the antagonistic pleiotropy hypothesis (AP); and
- the mutation accumulation hypothesis (MA).

Each hypothesis is coupled with a simulation component that can be switched 'on' or 'off' for any given simulation. The diversity hypothesis requires the existence of a disease population, which will exploit genetic relatedness. The antagonistic pleiotropy and mutation accumulation hypotheses require different chromosome components to be enabled, as discussed in Section 4 – namely, pleiotropic genes or mutation accumulation list respectively. It is expected that if any given component is disabled, then the associated hypothesis will not be active, either because the genetics are not present, with the non-adaptive hypotheses, or the benefit is absent, with the adaptive hypothesis. To conduct a complete comparative study of these hypotheses, various simulations were conducted with all the different combinations of these components enabled (see Table 1).

Figure 6(a) plots the evolution of the genetic expiry age gene for each of the experiments in Table 1. Figure 6(b) plots the corresponding proportion of deaths attributable to the gene. We expect that the genetic expiry gene will only have selective value in the simulations with the disease population enabled (DIV). As can be seen in Fig. 6(a), this is the case. In the

Table 1. Simulation experiment settings

Experiment	Diseases	Pleiotropic gene	Mutation accumulator
NULL	No	No	No
Diversity (DIV)	Yes	No	No
Antagonistic pleiotropy (AP)	No	Yes	No
Mutation accumulation (MA)	No	No	Yes
DIV & AP	Yes	Yes	No
DIV & MA	Yes	No	Yes
AP & MA	No	Yes	Yes
DIV & AP & MA	Yes	Yes	Yes

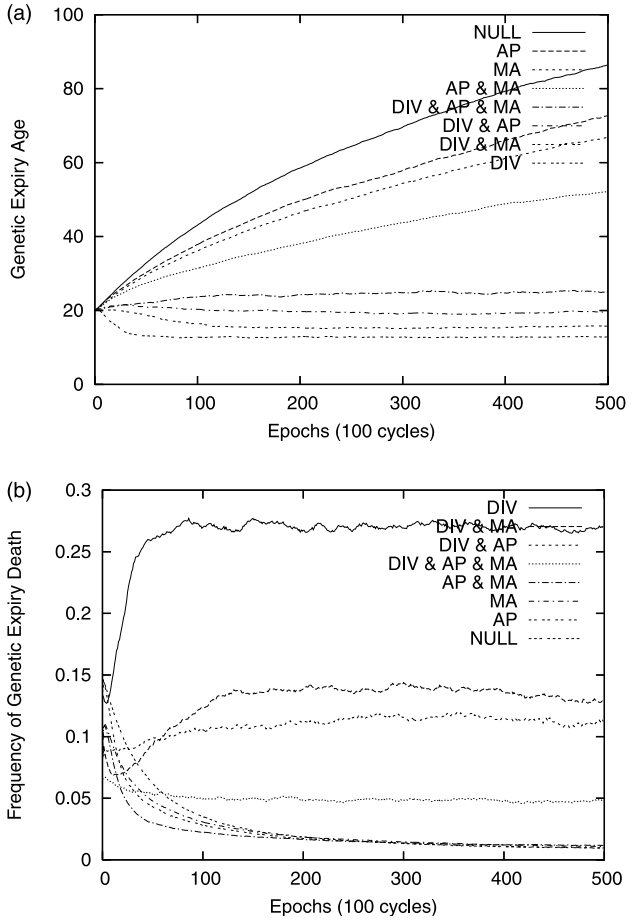

Fig. 6. Plots over experiments outlined in Table 1. Labels are sorted by final position. Panel (a) plots the evolution of the genetic expiry age gene that is present in all simulations. The evolution of ageing occurs only in the DIV simulations, with diseases enabled. Panel (b) plots the proportion of deaths caused by the genetic expiry age gene of panel (a).

remaining simulations, with diseases disabled (No DIV), Fig. 6(b) shows the genetic expiry age drops to background noise, below 1%. It was expected, in the simulations without diseases, that longer life spans would be selected for, as individual selection reigns supreme, and there could be no possible adaptive benefit to ageing. This can be seen in Fig. 6(a). The presence of the non-adaptive mechanisms (AP and MA), which also contribute to population turnover, cause the genetic expiry gene to evolve longer life spans. Hence, in Fig. 6(a), the DIV & AP and DIV & MA results are greater than DIV by itself, and the joint DIV & AP & MA is greater still.

Figure 7 shows the evolution of the non-adaptive genes (AP & MA). Figure 7(a) plots the evolution of the AP gene. This is the value that determines the ratio of food energy directed to fertility and health (maintenance), where higher values equate to a greater emphasis on maintenance and less on fertility. Figure 7(b) plots the evolution of the MA gene. To

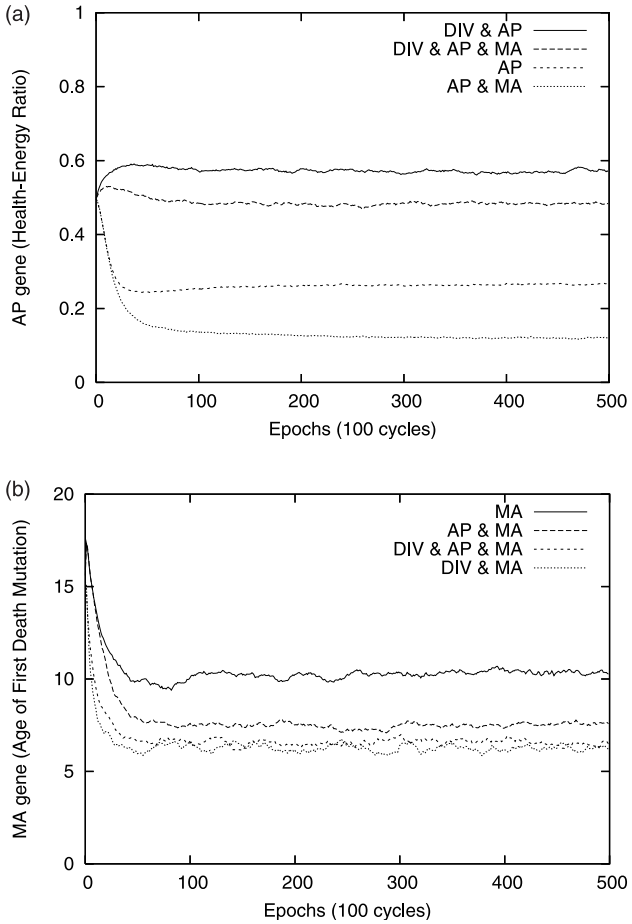

Fig. 7. Panel (a) plots, for the relevant runs, the antagonistic pleiotropy gene value, which determines the division of food to health and fertility stores – higher values signify less weight in health maintenance. Panel (b) plots, for the relevant runs, the age index of the first appearance of a death mutation in the mutation accumulation gene string.

simplify the representation of the MA gene string, this was measured as the first age appearance of a death mutation. Both the antagonistic pleiotropy and mutation accumulation hypotheses rely on a decreasing probability of reproduction with age (see Section 2), thus we would expect a synergy effect when the multiple ageing mechanisms are enabled (i.e. DIV or AP or MA). From the figures, this appears to be the case. For example, the AP and the MA genes, in Fig. 7(a) and 7(b) respectively, are selected towards greater ageing (i.e. lower AP and MA values) when both (AP & MA) components are enabled. When the disease component (DIV) is enabled, we would expect to see additional effects of the non-adaptive genes. First, as a consequence of the disease population, which parasitizes its requirements from the host health store, it is of greater importance for the host to maintain this health store. This would be expected to strongly factor in the AP trade-off. Figure 7(a) shows that when the disease (DIV) component is enabled, the AP gene evolves a value with a much greater weight in health maintenance. Second, as per our hypothesis, when the DIV component is enabled, a shorter life span has an adaptive benefit, and this would reduce the side-effect costs of the non-adaptive genes. This effect can be seen in both Fig. 7(a) and 7(b), where these runs evolve lower AP and MA genes. It is also noteworthy, in Fig. 7(b), that there is very little difference between the DIV & AP & MA run and the DIV & MA run, where the DIV component diminishes the synergy effect of the AP component on the MA gene.

As noted earlier, Williams (1957) outlined the major difficulties with an adaptive theory of ageing, one of which pertains to the infrequency of ageing deaths in natural populations. That is, a high rate of accidental death in nature would diminish the selection pressures for ageing. To address this criticism, all simulations were run with a high rate of accidental deaths (probability of 0.1 per cycle). Note that this was a required assumption for the antagonistic pleiotropy and mutation accumulation hypotheses. This did not cause a problem for our diversity hypothesis, which evolved shorter life spans regardless; there is still sufficient benefit in removing the lucky agents that survive to older ages. In fact, just as accidental deaths decrease the cost of ageing in antagonistic pleiotropy and mutation accumulation experiments, it will also decrease the individual cost of ageing in the adaptive experiments. This decrease in cost will balance the decrease in group benefit, perhaps making ageing more, but not entirely, selectively neutral.

6. CONCLUSION

As described in Section 2, Williams' main criticisms of Weismann's hypothesis, and more generally adaptive theories of ageing, were:

1. The fallacy of identifying senescence with mechanical wear.
2. The extreme rarity, in natural populations, of individuals that would be old enough to die of the postulated death mechanism.
3. The failure of gerontological research to uncover any death mechanism.
4. The difficulties involved in visualizing how such a feature could be produced by natural selection.

To conclude, we address each of these criticisms, except the third, which is outside the scope of this paper. (However, we note that this criticism applies equally to the alternative theories.)

The first of these criticisms relates particularly to Weismann's hypothesis of ageing. Apart from the increased probability of death with age, the host agents in our simulation retain a constant ability to reproduce throughout the duration of their lives. Thus we do not participate in Weismann's fallacy. Our diversity hypothesis, like Weismann's, holds that ageing is adaptive with the group benefit of making room for the young, although we formulate this as: populations that experience ageing turn over faster and therefore have greater diversity. To explore the second criticism, we included a high rate of accidental death in our simulations. The fourth criticism was formulated before the papers of Wynne-Edwards (1962) and Hamilton (1964) outlined selection mechanisms that favour altruistic adaptive explanations (see Section 3). Although the concept of group selection has remained contentious, inclusive fitness theory (kin selection) has gained much credibility since its formulation and provides a natural basis for a group selection mechanism.

In Section 3, we outlined an argument that group selection is supervenient on kin selection. The Price equation requires that for the group selection of altruistic genes, a positive association is necessary between group fitness and frequency of altruistic genes, which is likely to arise via kin selection. The reliance of a model on kin selection can be tested by the removal of kin associations – that is, ensuring that there is no correlation between the locations of the parent and child. Experiments have shown that such a removal of kin selection resulted in the collapse of group selection.

Our ALife simulations demonstrate that groups with shorter life spans do indeed have a greater diversity. This diversity has a fundamental effect in the co-evolutionary arms race between host agents and the disease agents that parasitize them. Host populations with shorter life spans, and therefore greater diversity, are less likely to experience a disease epidemic and can allocate more energy into sustaining a greater population of hosts. We implemented a group selection mechanism, where groups with larger host populations were more likely to produce successful migrant parties to found new groups and conducted a comparative study with two other popular non-adaptive explanations of ageing.

The evolution of ageing genes occurred readily in the simulations demonstrating the diversity hypothesis, although, when available, genes with other positive side-effects were selected over these simple genes.

In summary, our simulations show the viability of an adaptive account of accelerated ageing as a consequence of a combination of group and kin selection, incidentally demonstrating the utility of ALife simulation methods for developing and testing evolutionary hypotheses.

REFERENCES

Bredesen, D.E. 2004. The non-existent aging program: how does it work? *Aging Cell*, **3**: 255–259.

Edmunds, G.F. and Alstad, D.N. 1978. Coevolution in insect herbivores and conifers. *Science*, **199**: 941–945.

Epstein, J.M. and Axtell, R.A. 1996. *Growing Artificial Societies: Social Science from the Bottom Up*. Cambridge, MA: MIT Press.

Gilpin, M.E. 1975. *Group Selection in Predator–Prey Communities*. Princeton, NJ: Princeton University Press.

Goldsmith, T.C. 2004. Ageing as an evolved characteristic – Weismann's theory reconsidered. *Medical Hypotheses*, **62**: 304–308.

Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W., Railsback, S. *et al.* 2005. Pattern-oriented modeling of agent-based complex systems: lessons from ecology. *Science*, **310**: 987–991.

Hamilton, W.D. 1964. The genetical evolution of social behaviour, I and II. *J. Theor. Biol.*, **7**: 1–52.
Hamilton, W.D. 1975. Innate social aptitudes of man: an approach from evolutionary genetics. In *Biosocial Anthropology* (R. Fox, ed.), pp. 133–155. New York: Wiley.
Kim, J. 1993. *Supervenience and Mind*. Cambridge: Cambridge University Press.
Leroi, A.M., Chippindale, A.K. and Rose, M.R. 1994. Long-term laboratory evolution of a genetic life-history trade-off in *Drosophila melanogaster*. *Evolution*, **48**: 1244–1257.
Lewontin, R.C. and Dunn, L.C. 1960. The evolutionary dynamics of a polymorphism in the house mouse. *Genetics*, **45**: 705–722.
Mascaro, S., Korb, K.B. and Nicholson, A.E. 2001. Suicide as an evolutionarily stable strategy. In *Proceedings of the 6th European Conference on Advances in Artificial Life*, pp. 358–361, Prague, Czech Republic.
Maynard Smith, J. 1976. Group selection. *Quart. Rev. Biol.*, **51**: 277–283.
Medawar, P.B. 1952. *An Unsolved Problem in Biology*. London: H.K. Lewis.
Mitteldorf, J. 2004. Aging selected for its own sake. *Evol. Ecol. Res.*, **7**: 1–17.
Mitteldorf, J. 2006. Chaotic population dynamics and the evolution of aging. *Evol. Ecol. Res.*, **3**: 561–574.
Price, G.R. 1970. Selection and covariance. *Nature*, **227**: 520–521.
Skulachev, V.P. 1997. Ageing is a specific biological function rather than the result of a disorder in complex living systems: biochemical evidence in support of Weismann's hypothesis. *Biochemistry (Moscow)*, **62**: 1191–1195.
Weindruch, R. and Walford, R. 1986. *The Retardation of Aging and Disease by Dietary Restriction*. Springfield, IL: Thomas.
Weismann, A. 1889. *Essays Upon Heredity and Kindred Biological Problems*. Oxford: Clarendon Press.
Williams, G. 1957. Pleiotropy, natural selection, and the evolution of senescence. *Evolution*, **11**: 398–411.
Woodberry, O., Korb, K.B. and Nicholson, A.E. 2005. The evolution of aging. In *Proceedings of the Second Australian Conference on Artificial Life (ACAL 2005)*, Sydney, Australia.
Wynne-Edwards, V.C. 1962. *Animal Dispersion in Relation to Social Behavior*. Edinburgh: Oliver & Boyd.