

Environmental regulation can arise under minimal assumptions

J. McDonald-Gibson*, J.G. Dyke, E.A. Di Paolo, I.R. Harvey

Evolutionary and Adaptive Systems Group, Centre for Computational Neuroscience and Robotics, University of Sussex, Brighton BN1 9QH, UK

Received 25 August 2007; received in revised form 20 December 2007; accepted 20 December 2007

Available online 28 December 2007

Abstract

Models that demonstrate environmental regulation as a consequence of organism and environment coupling all require a number of core assumptions. Many previous models, such as Daisyworld, require that certain environment-altering traits have a selective advantage when those traits also contribute towards global regulation. We present a model that results in the regulation of a global environmental resource through niche construction without employing this and other common assumptions. There is no predetermined environmental optimum towards which regulation should proceed assumed or coded into the model. Nevertheless, polymorphic stable states that resist perturbation emerge from the simulated co-evolution of organisms and environment. In any single simulation a series of different stable states are realised, punctuated by rapid transitions. Regulation is achieved through two main subpopulations that are adapted to slightly different resource values, which force the environmental resource in opposing directions. This maintains the resource within a comparatively narrow band over a wide range of external perturbations. Population driven oscillations in the resource appear to be instrumental in protecting the regulation against mutations that would otherwise destroy it. Sensitivity analysis shows that the regulation is robust to mutation and to a wide range of parameter settings. Given the minimal assumptions employed, the results could reveal a mechanism capable of environmental regulation through the by-products of organisms.

© 2007 Elsevier Ltd. All rights reserved.

Keywords: Co-evolution; Gaia; Daisyworld; Niche construction; Rein control

1. Introduction

It is now widely accepted that organisms and their environment co-evolve. By their very nature, organisms affect their habitat through the consumption, transformation and excretion of resources. The scale of these effects may differ significantly in space and time. For example, burrowing earthworms change the composition of local soil (Lee, 1985), whereas the amplification of silicate rock weathering by plant life has resulted in the reduction in the atmospheric concentration of CO₂ by 10–100 times. Without the latter effect, the surface temperature of the Earth would be much higher, perhaps in excess of 50 °C (Schwartzman and Volk, 1989). The effects that organisms have on their environment, which have been referred to as ‘niche construction’ (Laland et al., 1996), can modify selective pressures and influence the course of evolution.

For example, niche construction can maintain or destroy stable polymorphisms, allow otherwise deleterious alleles to reach fixation, introduce an effect of evolutionary momentum, and influence dis-equilibrium (Laland et al., 1996, 1999; Silver and Di Paolo, 2006). Niche construction models have focused on the consequences of the organism/environment coupling for the evolution of a population. Investigation into environmental regulation as a potential consequence of this coupling has centred on the Gaia theory.

Gaia theory proposes that the system consisting of life, the oceans, atmosphere, and materials on the surface of the Earth is self-regulating, maintaining certain environmental variables within the limits necessary to sustain widespread life (Lovelock, 1988). The theory has remained somewhat controversial and has been the subject of a pointed critique. Some thought the theory teleological, implying foresight and planning on the part of the biota in order to maintain regulation. The debate is centred, however, on why organisms that are the product of Darwinian evolution,

*Corresponding author. Tel.: +44 1273872948.

E-mail address: jamie.mcdg@googlemail.com (J. McDonald-Gibson).

a process that focuses on immediate individual advantage, should act in ways that are beneficial to the biota (Dawkins, 1986).

Daisyworld (Watson and Lovelock, 1983) is a mathematical model designed to address some of these concerns. The name ‘Daisyworld’ refers to an imaginary planet that is warmed by a sun, and populated by black and white daisies that together regulate global temperature. The model is usually presented showing regulation around the optimal temperature for daisy growth despite increasing solar luminosity. Both the black and white daisies have an identical temperature-dependent growth function. This takes the form of an upside-down parabola that peaks at the optimum temperature for growth and declines on either side towards zero. Due to their lower albedo, black daisies absorb more solar energy than either white daisies or bare ground. In a cool environment, the solar energy that black daisies absorb causes their local temperature to rise. This increases their growth rate and generates a feedback loop with their rising numbers further warming the environment. This positive feedback on black daisy growth continues until the global temperature overshoots and subsequently stabilises close to the value that produces the highest daisy growth rate. Negative feedback constrains any further growth of black daisies, as further increases in temperature would now decrease their growth rate. The white daisies have a similarly regulatory effect and selective advantage in higher temperatures through cooling their environment. Environmental regulation in Daisyworld is achieved without recourse to teleology by the varying proportional coverage of black and white daisies.

There have been a number of modifications and extensions to the original Daisyworld model (see Wood et al., 2008 for a review). These and other models typically show regulation of environmental variables, such as temperature, through the interaction of individual species with each other and their environment. However, the regulation in each of these models requires a number of core assumptions. Lovelock (1988) and Lenton (2004) have both pointed out that the minimal requirements for global regulation may be life with its properties of growth and environmental alterations (A1), and bounds of habitability for life (A2). However, most mathematical models, simulations, and many explanations, invoke some of the following extra assumptions. Gaia theory models often presuppose niche construction traits that provide a selective advantage or disadvantage to the individual (A3). Often when such a trait has a selective advantage in a certain environment, it also improves environmental conditions for the whole population (A4). This is often implemented by allowing niche construction traits to have a stronger local than global effect (A5). In Daisyworld, for example, each daisy effectively maintains a distinct local environment that is crucial to the functioning of the model (Saunders, 1994). Most Gaia theory models also feature a single fixed environmental state at which maximum fitness (or growth rate) is obtained. For example, in Daisyworld

this takes the form of a temperature that gives maximum daisy growth; whereas in Downing and Zvirinsky’s (1999) simulation of regulation in chemical cycling networks, fitness is measured in relation to a preset ratio of chemicals. This assumption also employed in Daisyworld models in which the daisies can adapt towards the prevailing temperature such as Lenton and Lovelock (2000) and Wood et al. (2006). Both models use a similar growth function, which encodes a penalty for adaptation away from a predefined central value. In models that do allow the adaptation of organismic traits such as these and (Williams and Noble, 2005), constraints on adaptation, for example, limits on the rate of adaptation, are necessary in order for regulation to be maintained (A7). These assumptions are listed in Table 1.

In Daisyworld and many of its variants, assumptions A3–A6 combine to ensure regulation: local niche construction that benefits the individual also moves the global environment towards the state at which maximum fitness is obtained. The use of these assumptions has attracted criticism (Kirchner, 2003; Volk, 2004). Perhaps the most pertinent is that this mechanism, whereby certain traits have a selective advantage when they also contribute to regulation, has yet to be shown operating on a global scale.

This has been a factor in the recent shift in the discussion towards the possibility of global regulation arising through incidental by-products (Lenton and Wilkinson, 2003; Staley, 2002; Volk, 1998, 2004). Volk (2004) refers to Gaia as “life in a wasteworld of by-products” as exemplified by the atmosphere, which is a repository for wastes produced by metabolic processes where 99% of carbon dioxide, nitrogen, and methane, have been expelled by respirers, denitrifiers, and living prokaryotes, respectively. Organisms also consume and transform environmental materials. As such, the term ‘by-products’ includes incidental effects such as the amplification of silicate weathering by plants. If we assume that the environmental alterations responsible for regulation are by-products then it is necessary to dispense with assumptions A3, A4, and in many instances A5. This can result in a system that is robust to ‘cheats’ as by-products are incidental or obligate and consequently do not carry any fitness or energy cost (Wilkinson, 1999). By the same token, however, Kirchner (2002) argues that a trait yielding by-products beneficial for regulation will not necessarily propagate due to it benefiting carriers and non-carriers equally. Therefore, both environmentally beneficial

Table 1
Assumptions commonly required for regulation

A1	Organisms have an effect on their environment
A2	Organisms can only survive in a certain environments
A3	Niche construction traits provide a selective advantage or disadvantage
A4	Advantageous traits have an advantageous effect on the environment
A5	The effects of niche construction are stronger locally than globally
A6	There is a single fixed environmental state that gives maximum fitness
A7	There are constraints or limits on adaptation

and detrimental by-products are equally likely to spread. Critics of Gaia theory hold that any environmental regulation that does arise does so by chance, and that life generally just perturbs the environment subject to limits at the extremes (Volk, 2004). Whether there is an inherent tendency towards regulation in systems with widespread life or not, the most plausible candidate for the creation of any system of global regulation is some kind of mechanism acting on by-products (Staley, 2002; Volk, 1998, 2004). However, aside from environmental variables maintained in limiting states by negative feedback, a mechanism for global regulation through by-products still remains to be found (Lenton and Wilkinson, 2003).

In the following section, we present a model that invokes significantly fewer assumptions than previous models that have demonstrated environmental regulation. Models that have previously made progress towards reducing assumptions include Williams and Noble (2005) in which A6 was not present, and Staley (2002) which dispenses with assumptions A3, A4, A5, A7, while retaining A6. Staley, which features a different feedback structure to Daisyworld, only considers a single niche construction effect by allowing only white daisies. Daisyworld models that only allow a single daisy type and therefore a single environmental effect (and those that fix certain combinations of different optimal growth temperatures for each daisy type) can also reduce the number assumptions employed. However, Wood et al. (2008) point out that a more complete approach is to allow individuals to achieve a range of possible effects on the environment, as well as allowing adaptation towards a range of environments. Here, an environmental resource and a population of individuals that both adapt to and alter their environment is modelled. We allow individuals to evolve that either increase or decrease the environmental resource. In this respect, our model is similar to those Daisyworld models that feature both black and white daisies, and allow the optimal growth temperature for daisies to adapt towards the local temperature (Robertson and Robinson, 1998; Lenton and Lovelock, 2000; Williams and Noble, 2005; Wood et al., 2006).

The model we present features assumptions A1 and A2: all organisms modify their environment and can only survive in certain environments. In the model, an individual's fitness is based solely upon how well adapted it is to the shared global environment; there are no local environments. An individual's niche construction will have an effectively neutral effect on its fitness. Hence, assumptions A3, A4, A5 are not used. Furthermore, individuals are able to achieve equal fitness anywhere within the habitable range of resource values. Therefore, assumption A6 is not employed; there is no fixed point towards which regulation should proceed. We do include assumption A7 as follows: we assume that the rate of evolutionary adaptation must be constrained relative to the rate at which the niche construction activities of the population can change the resource levels. We consider the rate of change of adaptation to be a function of both the rate of

mutation, and the rate of the change in the frequency of different alleles occurring in the population due to reproduction. While adaptation typically takes many generations, environmental change through niche construction can affect the fit between organism and environment over faster time scales (Odling-Smee et al., 2003). We assume that life perturbs the environment significantly, and although organisms track environmental change through evolutionary adaptation, they cannot always immediately and exactly follow this change as fast as it occurs.

2. Methods

The simulation model we present shares a number of similarities with both Daisyworld and the niche construction models of Laland et al. (1996, 1999). A population of two-locus organisms that interact with their environment by increasing or decreasing an environmental resource is modelled. This resource determines the fitness (viability) of the organisms, which can be adapted to different resource levels. A real world example of such a resource could be pH, the concentration of a particular chemical in a well-mixed solution, the partial pressure of a gas, or temperature. External perturbation in the form of forcing can be applied to the resource.

The environment is represented by a single, global resource variable $R \in [0, 100]$. Individuals can survive in the environment when R is between 15 and 85; these are the bounds of habitability.

Each individual has a genotype that determines its effect on the resource and the resource-level to which it is best adapted. These traits are encoded in two loci, \mathbf{E} and \mathbf{A} , respectively, each containing a real-valued number between 0 and 1. The value at \mathbf{E} can be used to model a continuous range of niche construction effects. Here, \mathbf{E} is considered to take one of two possible alleles: e which reduces the resource, and E which increases the resource. Allowing just two opposite niche construction effects, as opposed to a continuous range, simplifies analysis and we have verified that it makes little difference to the results. An individual will have the e allele if $0 \leq \mathbf{E} < 0.5$ and the E allele if $0.5 \leq \mathbf{E} < 1$. The value at \mathbf{A} is scaled to specify the point that the individual is best adapted to within the habitable range of resource values $A \in [15, 85]$. It is used to calculate an individual's fitness using an inverted parabolic function of the resource that peaks at the A value and declines sharply on both sides towards zero. This function is similar to the growth function used in the Daisyworld model. The fitness of an individual i is a function of R and λ , where λ provides a measure of the span of the parabola:

$$F_i = \max \begin{cases} 1 - \frac{(A_i - R)^2}{\lambda^2} \\ 0 \end{cases} \quad (1)$$

Natural selection is modelled using a genetic algorithm with a proportion of the population subject to death,

selection, and reproduction at each time step. A constant death rate of γ of the population per unit of time is applied. Therefore, on average, the lifespan of an individual is $1/\gamma$. This may model an annual, daily or even hourly cycle. For each death event, a tournament, where the victor is the individual with higher fitness, is held between two randomly selected members of the population. If both individuals have the same fitness then one is randomly chosen to be the loser. The loser of the tournament is considered dead and removed from the population.

Both a fixed population and a population of varying size are modelled here. In the fixed population model, the winner becomes a parent and produces a single child that replaces the loser. If both individuals have zero fitness then the loser is replaced with a new random genotype. In the variable population model, the expected number of offspring produced by the winner, O_w , is determined by the fitness of winner and the size of the current population. The density dependant component is given by the discrete time logistic growth equation:

$$O_w = F_w \times e^{(r(1-(n/K)))}, \quad (2)$$

where F_w is the fitness of the winner and n is the number of individuals in the population. We set $r=1$ giving smooth density-dependent growth up to a carrying capacity, K . A minimum population size is enforced in the variable population version. If the population size falls below 10, then new randomly generated individuals are added to the population until it contains 10 individuals once again.

In both models, the offspring produced are clones of the parent that are subject to mutation. Mutation occurs with a probability of μ_A and μ_E at the **E** and **A** loci, respectively. This is modelled by adding a number drawn from a Gaussian distribution of mean 0.0 and standard deviation 0.05 to the allele. Reflexion is used to maintain the alleles between 0 and 1: if a mutation would result in a value outside of these bounds, then the amount that the value would be outside of the bound is subtracted from, or added to, the upper or lower bound, respectively.

The rate of change of the resource is a function of the population's niche construction effects and the external forcing:

$$\frac{dR}{dt} = \alpha \sum_{i=1}^n I_i + \beta(P - R). \quad (3)$$

If an individual i has the resource decreasing e allele then I_i is set to -1 ; if it has the resource-increasing E allele then I_i is set to $+1$. Constant α determines the strength of the niche construction such that $\alpha \times I_i$ is the environmental effect of the i th individual. The external forcing, P , is the value towards which the environment is being forced at time t . This is set to increase linearly from 0 to 100 over the course of a simulation run. Constant β determines the strength of the forcing. In the absence of any niche construction, the environmental resource would follow the linear increase of the forcing parameter. Real world

examples of such an effect include thermal driving operating on an entire biosphere from a star, or increasing chemical input into a freshwater lake via a runoff. We consider such forcing to be an external perturbation. Internal perturbations, originating from within the population, will also be produced via the stochastic death and reproduction of individuals.

The results presented in this paper were obtained using a Java implementation of the simulation model.¹ At each iteration of the simulation, the forcing target P is updated and the current resource value given by Eq. (3) is then calculated using Euler's forward method. Subsequently, death and reproduction are performed as described above.

This model minimally captures the co-evolution of life and the environment in a simple system consisting of a resource variable and a population of individuals that both adapt to and alter their environment. An individual's fitness is determined purely by how well it is adapted to the environment, and this itself is a heritable characteristic subject to selection and mutation. Maximum fitness can be obtained anywhere within the habitable range of resource values if an individual is adapted to the actual resource value. As such, there are no predefined optimal resource levels for the population. The strength of an individual's impact on the shared global resource is given by the α constant. For large population sizes and smaller values of α , the **E** trait will be effectively neutral for an individual. This can model a situation whereby individuals affect the environment through by-products that immediately disperse into the global environment. For regulation to arise, it cannot begin with individual selective advantage, but must occur at the population level. In this model, there is equal potential for niche construction that has a negative or a positive effect on the population as a whole, and there is no explicit selection for either type.

3. Results

In the following results, unless otherwise stated, the maximum population size, K , is set to 1000. The population is initialised with 1000 randomly generated genotypes. λ is fixed at 5 for all individuals; this gives a range of 10 resource values in which an individual will have non-zero fitness. The strength of the external forcing, β , is set to 0.01. The mutation rates, μ_A and μ_E , are fixed at 0.1. The death rate, γ , is fixed at 0.005. Simulations continue for 100,000 units of time. Iterations of the simulation take place every 0.2 units of time. These settings result in an average of a single reproduction tournament producing one death and new addition to the population per simulation iteration.

¹A fully commented sample implementation of the simulation in Java is available for download from <http://www.informatics.sussex.ac.uk/users/inmanh/EnvReg/simulation.zip>.

3.1. Without niche construction

Simulations were performed in which the population had no effect on the environment ($\alpha = 0$). In this scenario, as R increases due to the external forcing, the population continually adapts to the changing resource levels. The average resource value that the population is adapted to, \bar{A} , quickly converges to R and increases linearly as the external forcing increases. Results are plotted in Fig. 1.

3.2. With niche construction

Simulations were performed in which α was set to 0.0005 to enable life to have an effect on the environment. The α value and the maximum population size K determine the maximum effect that the population can have on the resource, and hence the maximum amount of external perturbation/forcing that can be regulated against. Given a strength of forcing of $\beta = 0.01$, with $\alpha = 0.0005$ and $K = 1000$, the maximum difference between the resource R , and forcing target P , that can be maintained by the population is 50 resource units.

With niche construction, regulation is swiftly established and maintained despite the forcing and perturbations. If regulation does break down, then after a small period of rapid environmental change, it is established once more. Each simulation run produces different stable states with regulation around different resource values.

Fig. 2 shows a typical simulation run with two periods of regulation. In this example, initially R increases linearly with P . However, once the resource nears the habitable range, the population's niche construction rapidly increases the resource level and the population adapts to these changes. A system regulating the resource against the forcing emerges from $time = 12,000$ (Fig. 2(a)). Subse-

quently, regulation is maintained through the total niche construction output of the population working against the forcing as shown in Fig. 2(b). This is also evinced in Fig. 2(d), which shows the transition from a population consisting predominantly of individuals with the E allele who increase resource levels, to one dominated by e allele individuals who reduce resource levels. The first period of regulation finishes at around $time = 66,500$, after which the resource rapidly increases towards the forcing target (Fig. 2(a)) and population fitness drops sharply (Fig. 2(c)). A second period of regulation is swiftly established from $time = 69,000$.

In order to emphasise the variety of dynamics possible under the same parameters, in Fig. 3 we show a different simulation run with three main periods of regulation. During these stable periods, the composition of the population changes in such a way that the total niche construction output of the population will oppose the effect of the forcing (Fig. 3(b)).

In both runs, the proportions of the two niche construction alleles in the population do not drift as was seen in the example without niche construction (Fig. 1(b)). Instead polymorphisms are evident that counteract the external forcing.

3.3. Sensitivity analysis

The results in the following section are based upon mean measurements of the percentage time that the resource was regulated, the number of distinct periods of regulation, and the duration of these periods, obtained over 100 simulation runs. [For information on how these statistics were calculated, see Appendix A.] The results of the sensitivity analysis are shown in Table 2.

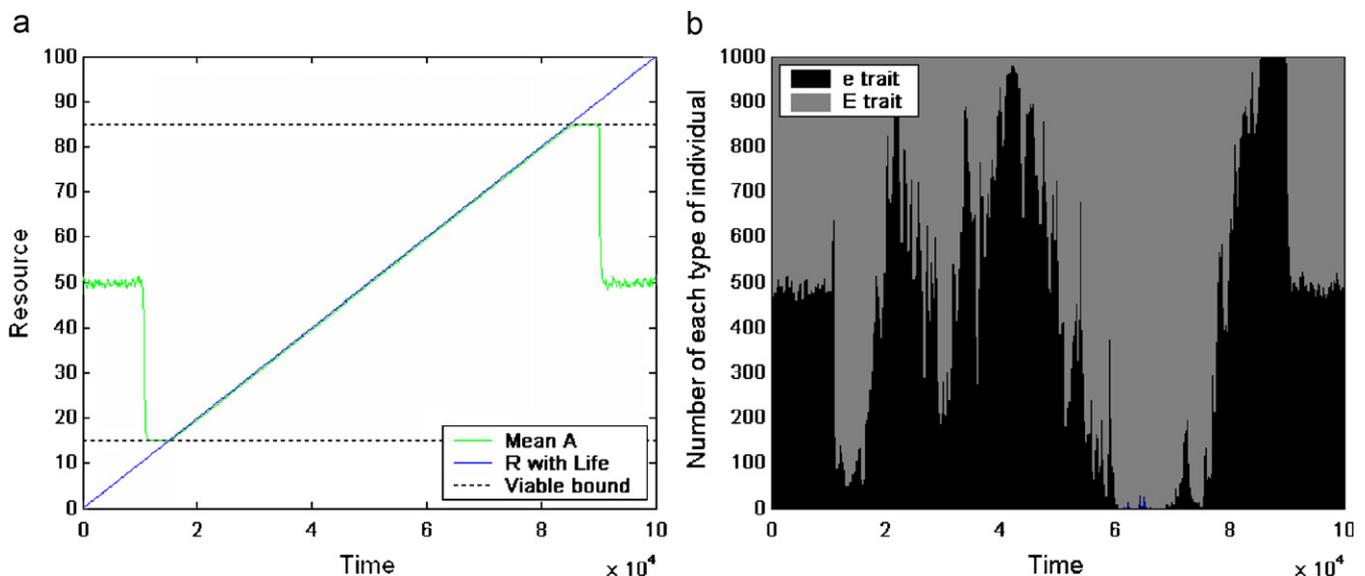


Fig. 1. Typical simulation run with no niche construction effect. Panel (a) plots the resource over time (blue line); mean \bar{A} is the mean resource level that the population is adapted to; the horizontal dotted lines are the bounds of habitability outside of which life cannot survive. Panel (b) is a stacked bar chart showing the number of individuals with the E allele and the number of individuals with the e allele.

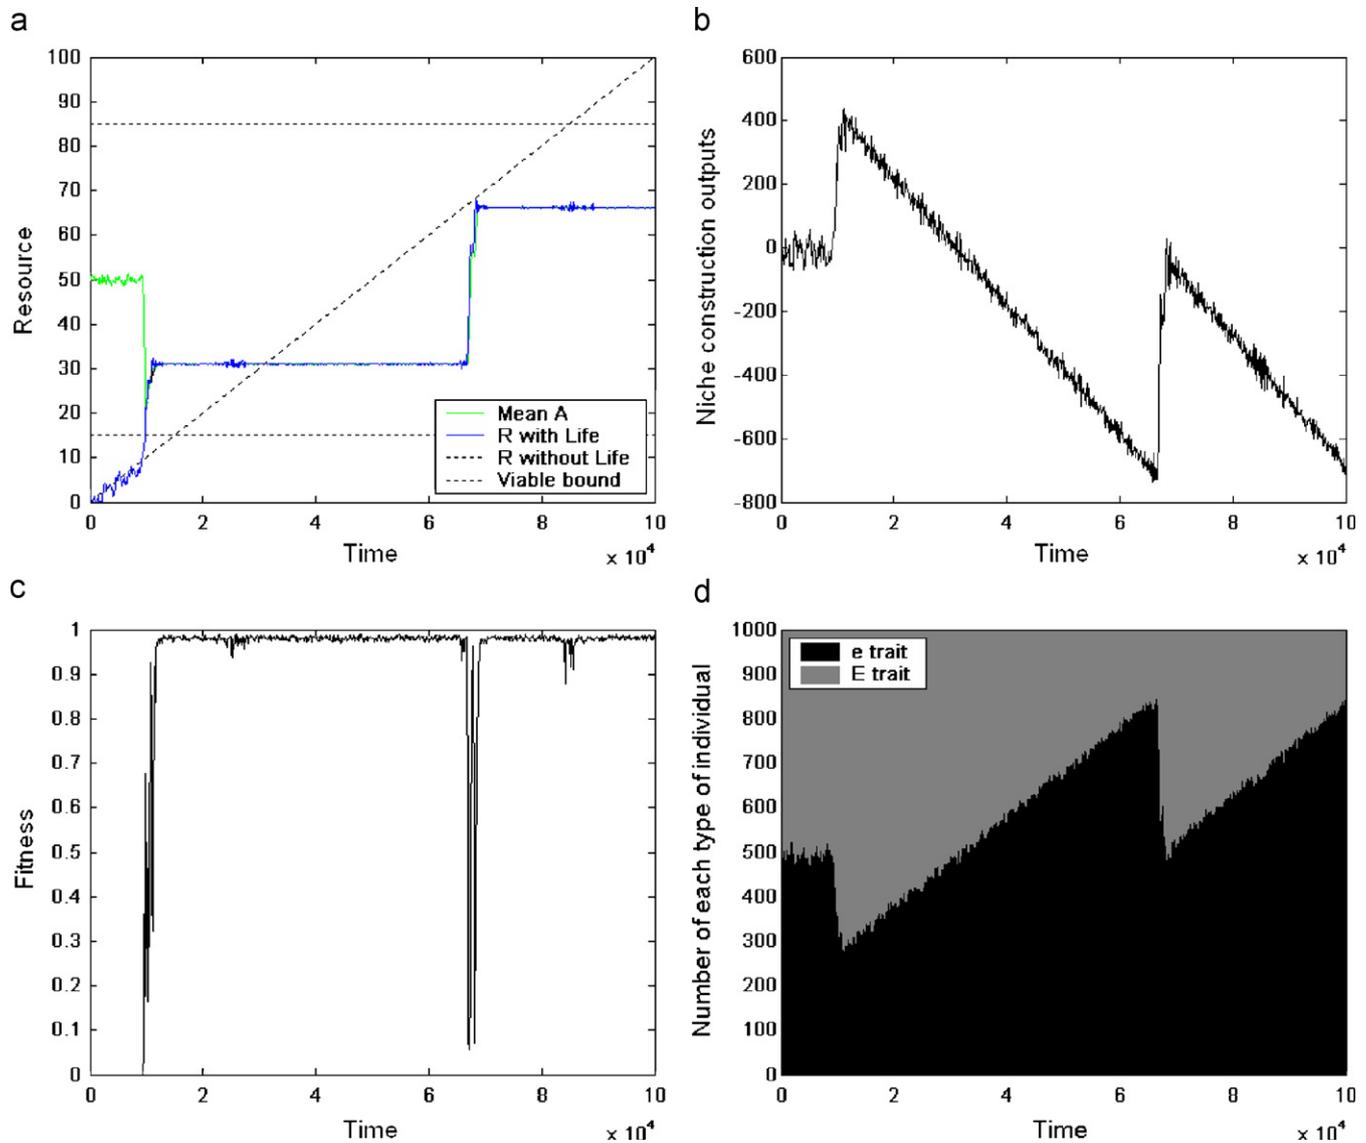


Fig. 2. Typical run of the simulation with two periods of regulation. Panel (a) plots the resource over time (blue line); mean A is the mean resource level that the population is adapted to; the diagonal dotted line is the forcing target P ; the horizontal dotted lines are the bounds of viability outside of which life cannot survive. Panel (b) plots the total niche construction output of the population. Panel (c) plots the mean fitness of the population. Panel (d) is a stacked bar chart showing the number of individuals with the E allele (who increase the resource), and the number of individuals with the e allele (who decrease the resource).

The version of the model with a variable population size governed by logistic regulates the resource for, on average, 85% of the time. This demonstrates that the regulation is not dependent on starting with the initial population size being equal to the carrying capacity K . Furthermore, the regulation does not depend on the size of the carrying capacity itself. Although we have for expediency used a relatively small carrying capacity to generate the results, larger population sizes with proportionally smaller individual effects (lower α values) display a similar level of regulation. As such, we believe the results here are applicable to a situation whereby the niche construction allele could represent a by-product that has an effectively neutral effect on an individual's fitness.

The mapping from the E locus to the e and E alleles can be changed to instead map to a continuum of niche construction effects of the over the range $[-1, 1]$. With this alteration, the total amount of regulation produced actually increases slightly (results are shown under the heading 'Continuous Impact Trait'). However, a greater number of distinct periods of regulation are observed.

The regulation is robust to changes in other parameters. The width of the fitness function, λ , can be freely adjusted without damaging regulation. The system also shows strong regulation for changes in the α value within one order of magnitude either side of the current value of 0.0005 (with β proportionally adjusted). Lower strengths of niche construction ($\alpha < 1.0 \times 10^{-5}$) cause the regulation to

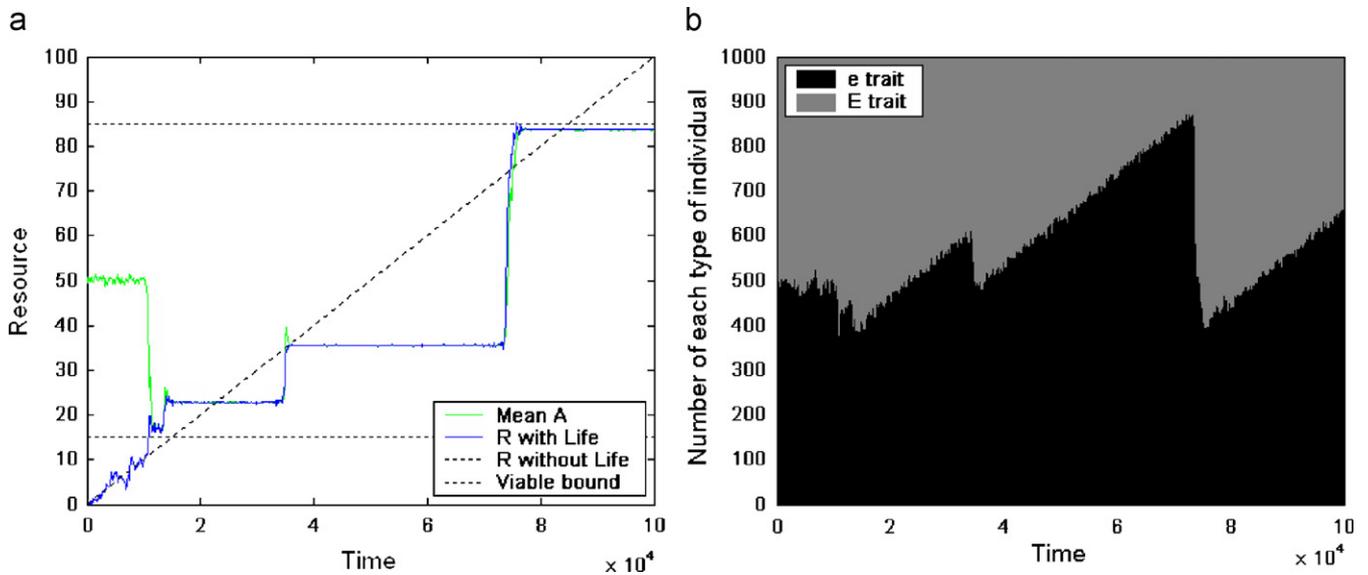


Fig. 3. Typical run of the simulation with three periods of regulation. Panel (a) plots the resource over time (blue line); mean A is the mean resource level that the population is adapted to; the diagonal dotted line is the forcing target P ; the horizontal dotted lines are the bounds of habitability outside of which life cannot survive. Panel (b) is a stacked bar chart showing the number of individuals with the E allele (who increase the resource), and the number of individuals with the e allele (who decrease the resource).

break down. This is perhaps surprising as, due to the proportional reduction in β , the population could still theoretically regulate against the same range of external perturbations. This is discussed in Section 4.1. While mutation does reduce the total amount of regulation, the regulation is robust to mutation rates of up to 0.3. This is also true when there are different mutation rates under this threshold at each locus. As mutation rates increase up to this threshold, the number of distinct periods of regulation rises and the length of these periods decreases. This indicates that mutation increases the likelihood of a given period of regulation breaking down. In populations with mutation rates the same at each locus, mutation rates of around 0.1 produce the most regulation (Table 2).

4. The mechanism of regulation

Regulation in the model can be explained in terms of niche construction and the maintenance of stable polymorphisms. Initial intuitions would suggest that as the external forcing increases, the population would just continually adapt to the rising resource levels. The regulation evidenced when the population exerts an effect on the resource emerges as follows: initially, the resource (R) begins below the viable range of any individual. As the forcing increases the resource level towards the lower limit of the habitable range, eventually an individual will have non-zero fitness. We call this individual A_{best} . Given a random initial population, it is equally likely to have an e or E allele. We assume that A_{best} has the resource-increasing E allele. As the frequency of A_{best} increases in the population, which it will do as it has the highest fitness, the frequency of the E allele increases, causing the total

impact of the population on the resource ($\sum_i I_i$) to increase. This moves R past the optimal resource value of A_{best} . The resource level will then increase towards the next individual with a higher A allele and so a new A_{best} will be established. Once again, this individual will have an e or E allele. If it too has the resource-increasing E allele then R will continue to increase, transit this new A_{best} , and move nearer towards the individual with the next highest A allele. Given a randomly generated initial population, it is likely that an individual with the resource-reducing e allele will be rapidly encountered. Once this happens, this type (which we will call $[e, A_{high}]$) will increase in frequency, decreasing $\sum_i I_i$, until it retards the resource back towards the previous A_{best} that has the E allele (which we will call $[E, A_{low}]$). This produces two dominant subpopulations, $[E, A_{low}]$ and $[e, A_{high}]$, which straddle the resource variable and force it in opposing directions as shown in Fig. 4. This process is subject to the effects of stochastic reproduction and death, as well as the perturbations that the rest of the population exert on the resource. Nevertheless, for a significant region in the parameter space, regulation towards a resource value, R^* , in-between the $[E, A_{low}]$ and $[e, A_{high}]$ types, is established in this way.

Once subpopulations of the $[E, A_{low}]$ and $[e, A_{high}]$ types have been established, regulation of the resource proceeds as follows: If the external perturbation P were to increase, R would increase and move closer to the optimum of the $[e, A_{high}]$ type. This would increase the frequency of e alleles relative to E alleles, lead to a decrease of $\sum_i I_i$, and so reverse the increase in R . If P were to decrease, the fitness of the $[E, A_{low}]$ type would increase, which would in turn lead to an increase in R once again. This results in R remaining regulated against varying P .

Table 2
Sensitivity of the regulation to parameter values and versions of the model

Model version/parameter settings	Percentage time in regulation (rounded)	No. of periods	Period length
Default	83	3.34	25,936
Logistic growth	85	5.73	18,233
Continuous impact trait	87	8.46	10,709
λ			
1	86	2.99	27,430
2	86	3.05	27,096
3	85	3.04	27,625
5	See default above		
10	74	3.77	21,293
15	75	5.44	16,466
20	82	5.12	14,755
K, α (averaged over 10 simulation runs)			
10,000, 5.0×10^{-5}	82	3.90	20,614
100,000, 5.0×10^{-6}	83	4.10	25,310
α, β			
0.05, 1.0	0	0	0
0.03, 0.6	39	5.06	7772
0.01, 0.2	71	3.52	20,431
0.005, 0.1	78	3.38	22,849
0.001, 0.02	79	3.56	21,587
0.0005, 0.01	See default above		
0.0001, 0.002	68	3.37	21,439
5.0×10^{-5} , 0.001	69	3.72	19,870
1.0×10^{-5} , 0.0002	29	2.94	10,221
5.0×10^{-6} , 0.0001	20	2.81	7785
μ_A, μ_E (both take the same value)			
0	83	2.00	38,415
0.01	70	1.54	41,714
0.03	75	2.09	35,163
0.05	78	2.40	32,301
0.08	81	3.09	27,509
0.1	See default above		
0.2	79	5.38	14,158
0.3	69	7.27	9171
0.4	33	7.12	1613
0.5	8	3.03	833
μ_A, μ_E (with different values)			
0.01, 0.1	84	3.29	25,624
0.05, 0.2	85	4.16	19,616
0.1, 0.01	78	1.87	41,325
0.2, 0.05	82	3.44	23,745

For each parameter value/model version, the table shows the mean percentage time that the resource was regulated, the mean number of distinct periods of regulation, and the mean length of these periods. These measurements were obtained over 100 simulation runs unless otherwise stated. Except for the parameters being tested, all of the other parameter values were unchanged from the default settings detailed at the beginning of Section 3. These statistics were taken during the period from $time = 15,000$ to $time = 85,000$ when the resource would be expected to be in the habitable range due to the external forcing.

Fig. 5 shows a close up view of the development of regulation in a single simulation run. The establishment of the $[E, A_{low}]$ and $[e, A_{high}]$ subpopulations is clearly evident in Fig. 5(a). During the formation of regulation, the size of the two subpopulations increases until the majority of the population belongs to one of these two groups (Fig. 5(b)). The resource plot in Fig. 5(a) shows that, for parameter values used, the regulation does not result in R remaining static between A_{low} and A_{high} . Instead, R oscillates

continuously around these values and R^* . The origin and impact of these oscillations is discussed in Section 4.3.

The balancing selective forces and the resulting polymorphic population of two dominant types that force the resource variable in opposing directions can be understood as an example of a *rein control system*. Rein control was introduced by Clynnes (1969) who hypothesised that certain variables (e.g., core body temperature in mammals) are maintained within a range of values by separate,

unidirectional control ‘reins’ that oppose forces that seek to perturb the variable. The rein control concept has since been developed in a physiological context (Saunders et al., 1998; Koeslag et al., 1999) and latterly applied to analysis of a simplified Daisyworld model (Harvey 2004; Dyke and Harvey, 2005, 2006). In this model, the $[E, A_{low}]$ and

$[e, A_{high}]$ subpopulations can be regarded as unidirectional control reins. Varying the strength of these reins (changing the numbers of individuals within each subpopulation) leads to the resource being regulated to within a narrow range over varying external perturbations.

4.1. The rates of adaptation and environmental change

The regulation described above arises in the simulation for a significant range of parameter values. Assumption A7, that the rate of adaptation must be constrained relative to the rate at which the niche construction activities of the population change the resource levels, is necessary. Both the rate of evolutionary adaptation via new mutations, and the rate of population turnover as existing subpopulations increase/decrease in size, are important.

For example, during the progression to regulation, when the $[E, A_{low}]$ subpopulation pushes the environment past its optima and towards the $[e, A_{high}]$ subpopulation, R must be changing fast enough relative to the change in proportions of the two subpopulations, such that when one subpopulation is fitter than the other, the fitter one does not take over the entire population at the expense of the temporarily less fit. Similarly, during regulation, changes in the resource level away from R^* must be corrected before the extinction of the subpopulation that is temporarily less viable. In either case, the rate of adaptation must be sufficiently constrained to prevent the fittest subpopulation from immediately adapting to the changing resource levels by fixing mutations at the A locus.

The α parameter, which determines the strength of the niche construction effects, is crucial in satisfying the above conditions. The sensitivity analysis shows that if α is set too

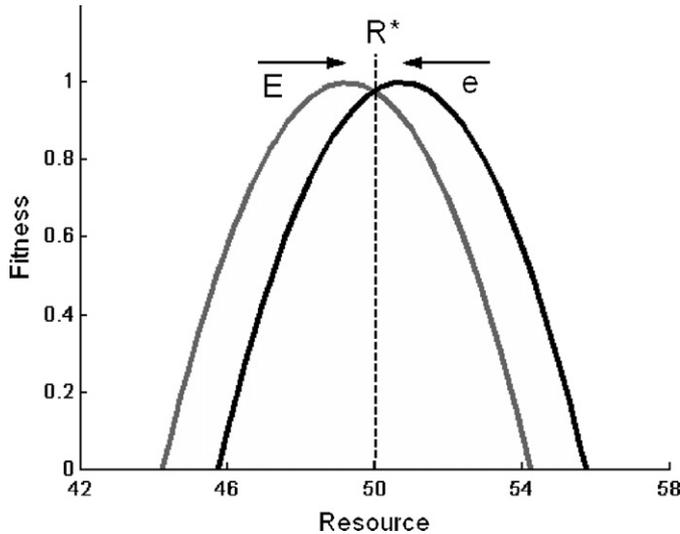


Fig. 4. The fitness range of the $[E, A_{low}]$ type is plotted with a solid grey line; the $[e, A_{high}]$ type's fitness range is plotted with a solid black line. The vertical dashed black line marks the value for R^* . The resource will be maintained, on average, around this value. When R is closer to the $[E, A_{low}]$ type, E alleles increase in frequency relative to e alleles, raising resource levels. Conversely, when R is closer to the $[e, A_{high}]$ type, the relative fitness of the two types leads to the resource decreasing. The niche construction of each subpopulation pushes the resource towards R^* , thus counteracting the niche construction of the opposing subpopulation, and together regulating R against perturbations.

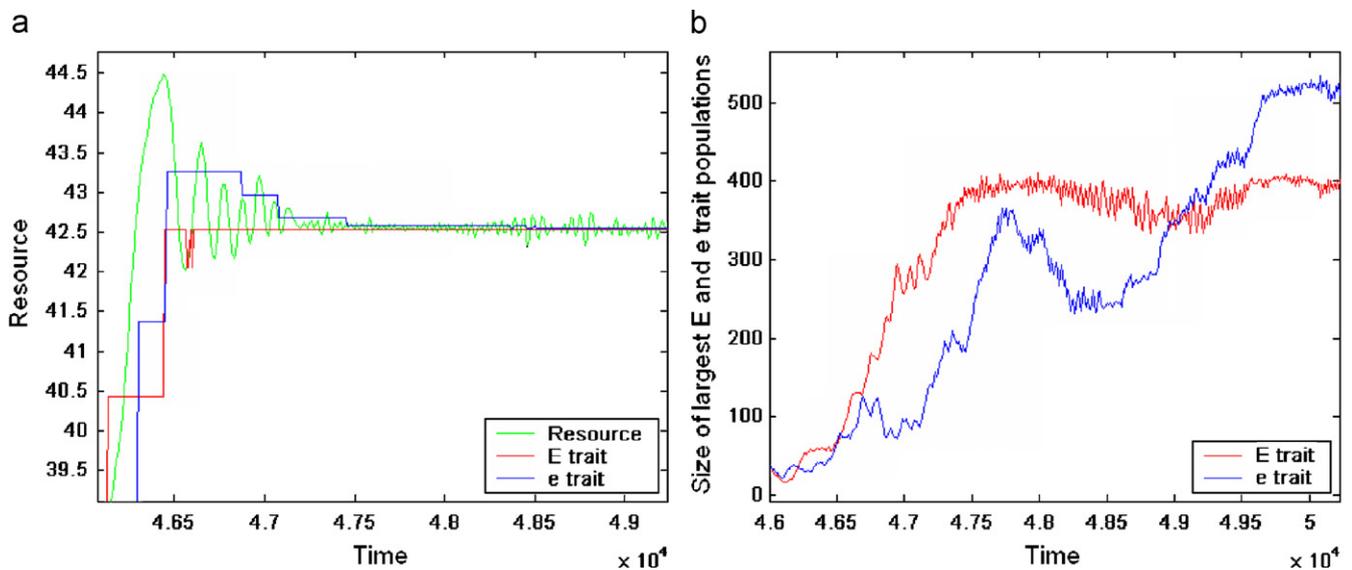


Fig. 5. Close up view of the establishment of regulation during a typical simulation run. Panel (a) shows the resource levels to which the largest E and e allele subpopulations are adapted; panel (b) shows the size of these two subpopulations. Panel (a) shows the resource (green line) becoming regulated around the two subpopulations, with $[E, A_{low}]$ and $[e, A_{high}]$ visible in red and blue respectively. During the establishment of regulation, the size of the two subpopulations increases until the majority of the population belong to one of these two groups (b). Note that even after regulation is established there are continuous oscillations in the resource.

low (less than 1.0×10^{-5} , with β set to 2.0×10^{-4}), then there is little regulation. This is likely to be because population effects are not strong enough to change the environment at a fast enough rate for assumption A7 to be satisfied.

The population can and does adapt to changes in the resource wrought by its own niche construction and the external forcing. However, the population driven changes in the resource can occur over a shorter timescale than the timescale over which evolutionary adaptation operates.

4.2. How mutation can break regulation

The sensitivity analysis showed that the regulation is robust to mutation rates of up to 30%. Without mutation, once regulation is established it is likely to continue until the external forcing becomes too strong for the population to counteract. Mutation at either the **E** or **A** loci can destroy regulation earlier than this.

We have postulated that regulation is maintained by two subpopulations acting as control reins: the $[E, A_{low}]$ subpopulation opposes decreases in R and the $[e, A_{high}]$ subpopulation opposes increases in R . Genetic drift in the **E** trait can lead to the gradual weakening of these control reins as E alleles mutate to e alleles and vice versa. This process can continue until one of the reins is not strong enough to counteract changes in the resource. At this point, internal or external perturbations will eject R from between A_{low} and A_{high} and the regulation will break down.

Mutation at the **A** locus can destroy the regulation as follows: Mutants from either the $[E, A_{low}]$ or $[e, A_{high}]$ subpopulations with an **A** allele between A_{low} and A_{high} will theoretically have a selective advantage over their parents as they will be closer to R^* . This results in the $[E, A_{low}]$ and

$[e, A_{high}]$ subpopulations creeping towards R^* . If this continues until A_{low} and A_{high} share the same **A** trait and hence have the same fitness, then regulation breaks down. This is because the frequency of either subpopulation, and hence either e or E alleles, can no longer increase at the expense of the other in response to changes in the resource. The model proves surprisingly resilient to such intermediate mutations at the **A** locus. What tends to happen in simulation is that the two subpopulations creep towards each other and then stabilise very close together, typically around 0.005 resource units apart. This process is evident in Fig. 5(a) above. Given a continuous number of **A** alleles that may be realised in any genotype, as $A_{high} - A_{low} \rightarrow 0$, the fitness gain that a mutant could achieve and the probability of an intermediate mutation both $\rightarrow 0$. Crucially, during regulation the resource constantly fluctuates outside the range spanned by A_{low} and A_{high} . This reduces the amount of time R spends within this range and consequently decreases the selective pressure for **A** alleles between A_{low} and A_{high} . Internal perturbations produced by stochastic fluctuations in the population can cause these changes in R . Furthermore, the population's niche construction leads to oscillations in the resource that produce a similar effect.

4.3. Population driven oscillations in the resource levels

Continuous oscillations in the resource around R^* which move the resource beyond the range spanned by the $[E, A_{low}]$ and $[e, A_{high}]$ types are observed in simulation and can be seen in Fig. 5(a). The oscillations arise due to the time lag in the population's response to the rapidly changing environment. This effect is shown in Fig. 6, which displays a close up view of a typical oscillation during regulation from a single simulation. This plot is

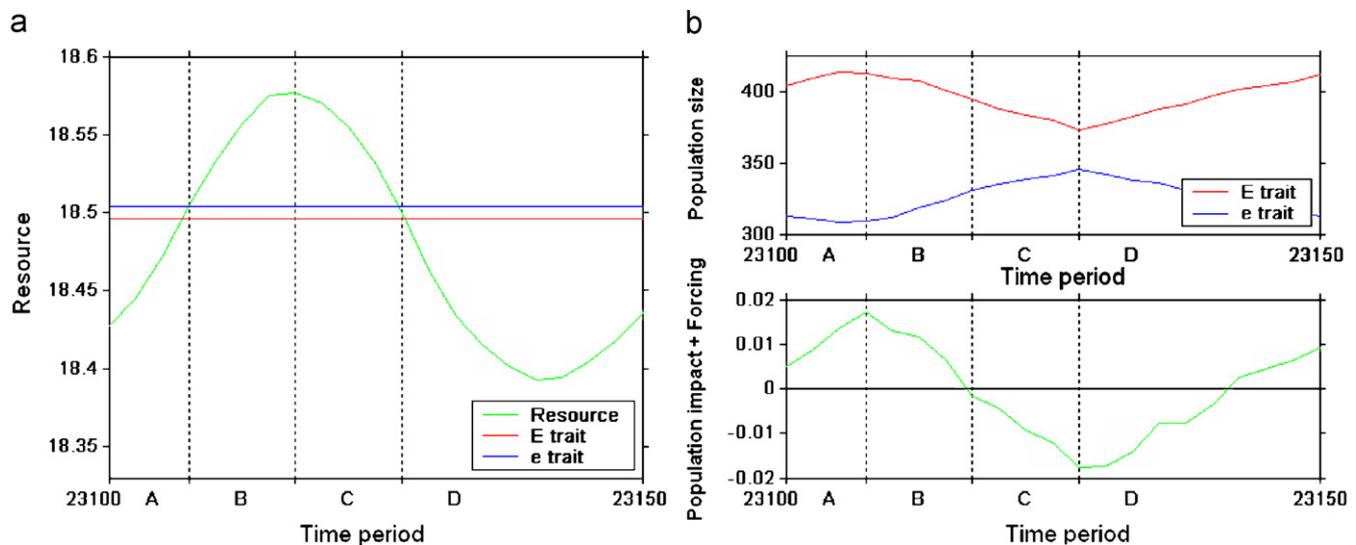


Fig. 6. Close up view of a typical oscillation in the resource. The vertical dotted lines separate the oscillation into four time periods: A, B, C, D. Panel (a) shows the resource (green line), and the resources levels that $[E, A_{low}]$ and $[e, A_{high}]$ are adapted to (the red and blue lines respectively). Panel (b) shows the size of the $[E, A_{low}]$ and $[e, A_{high}]$ subpopulations. It also shows total effect that the population and forcing has on the resource at each time step (green line). When this is negative the net effect of the niche construction and forcing decreases the resource; when it is positive the resource will be increased.

separated in to four time intervals: A, B, C, D. During period A, R is rising and is closer to the $[E, A_{low}]$ subpopulation. The $[E, A_{low}]$ type is therefore fitter than the $[e, A_{high}]$ type; consequently, the frequency of E alleles relative to e alleles is increasing. Note that the frequency of E alleles is increasing beyond the minimum amount needed to increase resource levels (Fig. 6(b), bottom). Once $R > R^*$ (period B), this relationship will be reversed and e alleles will begin to replace E alleles in the population (Fig. 6(b), top). However, the resource does not start decreasing immediately. This is because the ratio of E to e alleles is still high enough to keep R increasing (Fig. 6(b), bottom). The rate that the population adapts to the higher resource levels through the substitution of e alleles for E alleles is not fast enough to prevent the overshoot of R^* . At the start of period C this ratio of E to e alleles has become low enough to reverse the direction of the resource (this is shown by the first intersection of the plotted line with the x -axis in Fig. 6(b), bottom). Subsequently $[e, A_{high}]$ will continue to have a selective advantage as R moves back down towards R^* . During this period, the frequency of e alleles relative to E alleles increases beyond the minimum amount needed to reduce R (Fig. 6(b), bottom). This period of selective advantage for $[e, A_{high}]$ sets up the next overshoot of R^* at the start of period D and ensures the continuation of the oscillations.

This snapshot in Fig. 6 was taken in the middle of a period of regulation that lasted for 25,000 time steps. During this time, the average fitnesses of $[E, A_{low}]$, $[e, A_{high}]$, and at R^* , were 0.99822, 0.99801 and 0.99816, respectively. Had the resource remained at R^* and not oscillated, then any individual adapted to R^* would have had a fitness of 1, and individuals at A_{low} and A_{high} would have had a fitness of 0.96. Without oscillations, there would have been a greater selective pressure for subpopulations with A alleles in-between A_{low} and A_{high} .

The oscillations in the resource are driven by the changes in the proportions of the niche construction traits over very short timescales. These oscillations are likely to inhibit the fixation of certain mutations detrimental to regulation, which would otherwise occur over longer timescales in a non-oscillatory environment.

5. Discussion

The model presented here features a single global resource variable that can be regulated around contingent stable states via the niche construction activities of the population. The principal mechanism underlying regulation was shown to be two subpopulations operating as a rein control system: one subpopulation, adapted to a slightly lower resource level, increases the resource; the other, adapted to a slightly higher resource level, decreases the resource. Each subpopulation's niche construction counteracts certain perturbations and the niche construction of the other subpopulation. This can result in oscillations in the resource around these two subpopula-

tions because of the lag in the population dynamics relative to the quickly changing environment. These oscillations are likely to prolong the lifespan of a period of regulation by suppressing the fixation of mutants adapted to resource values in-between the optima of the two subpopulations. The sensitivity analysis showed that the regulation was robust to changes in parameter values and that allowing mutation at both loci did not significantly alter the regulation.

The individuals modelled here can both adapt to and alter their environment. In this respect, the model is similar to those Daisyworld models that allow the daisies to adapt towards their local temperature. Of these, the Daisyworlds of Robertson and Robinson (1998) and Wood et al. (2006) do not produce stable environmental states, although the latter does produce regulation in the form of temperature oscillations between the bounds of habitability. In common with Lenton and Lovelock (2000) and Williams and Noble (2005), we find that such adaptation need not prevent stable regulating states emerging. During a stable period, despite the genetic variation possible in the population, regulation is achieved by two dominant subpopulations, each of which regulates against perturbations to the resource in a single direction. Many earlier Daisyworld models that allow a greater range of daisy traits, for example, those that allow a finite number of possible daisy shades (Lovelock 1992; Lenton, 1998; Lenton and Lovelock, 2001), similarly produce two dominant types at any one time during regulation. Environmental oscillations have also been reported previously in Daisyworld models. Recently, Wood et al. (2006) demonstrated that long period environmental oscillations between the limits of habitability could arise as a consequence of the lag in a population's response to environmental change. Our results also show internally generated oscillations due to a similar lag. Here, the amplitude and period of the oscillations is much shorter as they are driven by the changes in the proportions of niche construction traits over faster timescales than adaptation operates over. Furthermore, the oscillations are likely to play an important role protecting the regulation against mutants that could otherwise destroy it. The benefit that niche construction induced oscillations may yield to a population has also been highlighted by Boyle and Lenton (2006), who showed that environmental fluctuations could limit the spread of parasites within groups of cooperating individuals.

The regulation evident in the results presented here requires very few assumptions in comparison to previous models. Daisyworld and many other models feature local environmental conditions that differ from the global conditions. This differentiation enables a mechanism whereby niche construction that benefits the individual also contributes to global regulation. We have shown that such an assumption is not a prerequisite for regulation: A resource variable can be maintained within certain limits via the changes in frequency of niche construction alleles that confer no immediate benefit to the individuals that carry them.

There has been speculation that a system of global regulation constructed from the by-products of organisms would be a chance occurrence (Kirchner, 2003; Volk, 2004). In these simulations, an inherent tendency towards regulation through traits that could model by-products was evident. Kirchner (2002, 2003) have argued that a trait yielding by-products beneficial for regulation would not necessarily spread due to it benefiting carriers and non-carriers equally. The perspective of considering who benefits from a given niche construction trait is not wholly appropriate here. The $[E, A_{low}]$ and $[e, A_{high}]$ subpopulations both push the environment away from their optimum resource levels. As such, neither benefits from their niche construction when they are considered in isolation. The system can only really be considered as a whole, the dynamics of which lead to two subpopulations, each counteracting the environmental change of the other, together enabling regulation.

Rather than regulate the resource around a predefined fixed point, the system passes through a different series of stable states during any particular simulation run. During the rapid transitions between periods of regulation, there are mass extinctions and expansions, before the system stabilises on an often qualitatively different resource level and population composition. This is consistent with the real world. The evolutionary history of species often reveals punctuated equilibria: periods of evolutionary stasis interrupted by rapid innovation and change (Eldredge and Gould, 1972). The evolution of ecosystems as evinced by the history of marine biota (Lenton et al., 2004), can also proceed with similar dynamics. Critics of Gaia theory have cited the biological amplification of some recent rapid changes in climate as contrary to notions of global regulation (Kirchner, 2002, 2003). Lenton and Wilkinson (2003) point out that while in the short term the biota may destabilise the system, from the perspective of Gaia theory, such perturbations tend to be corrected in the long run. In the results presented here, unstable periods dominated by positive feedback may be observed, but within the limits of habitability, these are quickly transitioned towards stable regulating states. Regulation emerges because it is very probable that opposing resource-increasing and resource-decreasing subpopulations will arise swiftly in the development of the system. The system keeps quickly changing until it falls into such an attractor. In this respect, its behaviour is reminiscent of Ashby's cybernetic model of adaptation, the ultrastable system (Ashby, 1960). The Earth and its biota can be considered as a dynamical system that may visit numerous attractor states as it progresses along a particular co-evolutionary trajectory. Internal or external perturbations may result in new stable states being found. The simple model presented here exhibited this behaviour. We argue that such a perspective, rather than an assumption of regulation around fixed points, is more relevant when considering regulation in systems that are comprised of organism and environment couplings.

Following the collapse of a regulating regime, a new period of regulation appeared more likely to commence if there were certain rates of mutation. In systems with mutation, mutation rates of around 0.1 produced the most regulation. These results could suggest that some diversity in the population could increase the probability of stable states arising. Dyke et al. (2007) suggested that if a population contains a variety of different genotypes, perhaps through mutation, then each of these could be instrumental in seeding future periods of stability. The genetic variation maintained could therefore enhance the 'evolvability' of the system and its components, affecting long-term evolutionary dynamics. Future work could further investigate the diversity–stability relationships in systems regulating in the manner presented here.

The development and maintenance of regulation was a key factor in determining the composition of the population, including the characteristics of the inhabitants. Laland et al. (1996, 1999) have used two-locus population genetic models to study the evolutionary consequences of niche construction. They found that niche construction can support stable polymorphisms where none are expected. In the results herein, polymorphisms were clearly evident in the proportions of niche construction traits. These traits did not drift as could usually be expected of selectively neutral traits, but were instead maintained and indeed tuned in such a way that resulted in environmental regulation. In the real world, a correlation between such polymorphisms and an environmental variable could provide a signifier for this form of regulation. Another potential manifestation is the characteristic oscillation around the stable point. It is worth reiterating that in order for such regulation to occur, the niche-construction traits must be selectively neutral, probably incidental or obligate by-products. Furthermore, the rate of adaptation must be constrained relative to the rate at which the niche construction activities of the population can change the resource levels. Systems that may satisfy these criteria range from highly localised microbial ecosystems through to organism and environment couplings that affect the global environment. Environmental characteristics that are potential candidates for regulation include temperature, salinity, and pH, as well as the concentrations of various substances in the soils, atmosphere and oceans.

Increasing the complexity and realism of the model with, for example, the incorporation of factors such as predation, cyclical extinctions, and thermodynamic constraints, would provide an important test of the regulatory mechanism. It would also be worthwhile to model a situation whereby the organisms that affect the environmental resource are not the ones whose fitness is directly contingent on its value. Regulation could still arise if the niche constructor is linked to other organisms dependent on the resource through predator/prey relationships or intermediate abiotic systems. Odling-Smee et al. (2003) have highlighted the importance of such links, terming them 'environmentally mediated genotype associations'

(EMGAs). Future work could make the model specific to candidate real-world environmental variables through the addition of the pertinent EMGA relationships. Additionally, mathematical models could be employed to investigate aspects of the regulation highlighted by the simulation results. Quantifying exactly how much the oscillations in the resource suppress the intermediate A trait mutations which damage regulation would be a good candidate for this type of investigation. It would also be worthwhile to further specify the rate at which the environment has to change relative to the rate of adaptation for regulation to be observed.

The interactions between the resource-increasing and resource-decreasing subpopulations shown here somewhat resemble producer/consumer relationships. Such relations are widespread in nature and various authors have highlighted their importance to Gaia theory (Volk, 2004). For example, organisms can evolve to consume excess pollutants that had been building up to the detriment of the system. This can enable systems of chemical or nutrient cycling to emerge. Downing and Zvirinsky's (1999) 'Guild' simulation explores environmental regulation in chemical cycling networks, demonstrating regulation towards an optimum concentration of chemicals. A follow up called 'Metamic' (Downing, 2002) found that, with the incorporation of energy costs for niche construction, regulation became much less probable. Both simulations employed similar assumptions to Daisyworld. In a similar vein to the work presented here, Williams and Lenton's (2007) 'Flask' model of nutrient recycling ecosystems relaxed the assumptions employed by Guild and Metamic. They found that efficient cycling networks could still emerge, but so-called 'rebel' organisms that exploit under-utilised resources could lead to mass or even complete extinctions of the population. A further line of investigation would be to examine the possibility of environmental regulation in cycling networks arising through the specific mechanism presented here.

We believe that this model addresses a number of criticisms directed towards the Daisyworld model and Gaia theory more generally. In particular, regulation does not have to require predefined optimal resource levels, nor differing local environments. We have shown how Darwinian selection operating on organisms that have a significant impact on their environment can produce an efficient rein control system that is able to regulate a resource variable against a wide range of perturbations.

Acknowledgments

We thank Tim Lenton and an anonymous reviewer for their valuable comments on the manuscript.

Appendix A

The sensitivity analysis recorded the mean percentage time that the resource was regulated, the mean number of

distinct periods of regulation, and the mean length of these periods.

A period of regulation was identified by monitoring the mean resource value \tilde{R} and mean A trait value \tilde{A} (the resource value that the population is adapted to) over intervals 200 time units long. Mean measurements are necessary as during regulation the resource fluctuates and oscillates around the stable point. Regulation of the resource was identified when the following conditions were met over two or more consecutive intervals:

1. $|\tilde{R} - \tilde{R}_1| < 1$: The resource remains within a small range.
2. $|\tilde{A} - \tilde{A}_1| < 0.25$: The population is largely fixed around a single resource value.
3. $|\tilde{R} - \tilde{A}| < \lambda$: The population is adapted to the resource and is viable.

The subscript 1 above denotes mean readings from the first interval examined in a given period of regulation. To avoid false positives, a period of regulation was only recorded if there were more than eight sequential intervals fulfilling the criteria above. The beginning of regulation is recorded as being the start of the first interval identified. The end of regulation is recorded as the end of the last interval that fulfils the criteria above. This conservative measure of regulation means that the results are likely to slightly underestimate the total amount of regulation.

Some parameter values tested could result in the resource entering the habitable range of resource values $R \in [15, 85]$ earlier, or remaining for longer within this range, than others. To ensure a fair comparison, results were recorded only within the period from $time = 15,000$ to $time = 85,000$ when the resource would be expected to be in the habitable range due to the external forcing.

References

- Ashby, W.R., 1960. Design for a Brain: The Origin of Adaptive Behaviour, second ed. Chapman & Hall, London.
- Boyle, R.A., Lenton, T.M., 2006. Fluctuation in the physical environment as a mechanism for reinforcing evolutionary transitions. *J. Theor. Biol.* 242 (4), 832–843.
- Clynes, M., 1969. Cybernetic implications of rein control in perceptual and conceptual organization. *Ann. N. Y. Acad. Sci.* 156, 629–670.
- Dawkins, R., 1986. *The Blind Watchmaker*. Longman.
- Downing, K., 2002. The simulated emergence of distributed environmental control in evolving microcosms. *Artif. Life* 6 (2), 123–153.
- Downing, K., Zvirinsky, P., 1999. The simulated evolution of biochemical guilds: reconciling Gaia theory and natural selection. *Artif. Life* 5 (4), 291–318.
- Dyke, J., Harvey, I.R., 2005. Hysteresis and the limits of homeostasis: from Daisyworld to phototaxis. In: Caparrere, M., Freitas, A., Bentley, J., Johnson, C., Timmis, J. (Eds.), *Proceedings of VIIIth European Conference on Artificial Life, ECAL 2005*, Springer, pp. 332–342.
- Dyke, J., Harvey, I.R., 2006. Pushing up the daisies. *Artificial Life X*. In: Rocha, L.M., Yager, L.S., Bedau, M.A., Floreano, D., Goldstone, R.L., Vespignani, A. (Eds.), *Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, MIT Press, pp. 426–431.

- Dyke, J.G., McDonald-Gibson, J., Di Paolo, E., Harvey, I.R., 2007. Increasing complexity increases stability in a self regulating ecosystem. To appear in Proceedings of IXth European Conference on Artificial Life, ECAL 2007, Lisbon, Portugal.
- Eldredge, N., Gould, S.J., 1972. Punctuated equilibria: an alternative to phyletic gradualism. In: Schopf, T.J.M. (Ed.), *Models in Paleobiology*. Freeman Cooper, San Francisco, pp. 82–115.
- Harvey, I.R., 2004. Homeostasis and rein control: from Daisyworld to active perception. In: Pollack, J., Bedau, M., Husbands, P., Ikegami, T., Watson, R.A. (Eds.), *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems, ALIFE'9*, pp. 309–314.
- Kirchner, J.W., 2002. The Gaia hypothesis: fact, theory, and wishful thinking. *Climatic Change* 52, 391–408.
- Kirchner, J.W., 2003. The Gaia hypothesis: conjectures and refutations. *Climatic Change* 58, 21–45.
- Koeslag, J., Saunders, P., Wessels, A., 1999. The chromograinins and the counter-regulatory hormones: do they make homeostatic sense? *J. Physiol.* 517, 643–649.
- Laland, K.N., Odling-Smee, J., Feldman, M.W., 1996. The evolutionary consequences of niche construction: a theoretical investigation using two-locus theory. *J. Evol. Biol.* 9, 293–316.
- Laland, K.N., Odling-Smee, F.J., Feldman, M.W., 1999. Evolutionary consequences of niche construction and their implications for ecology. *Proc. Natl. Acad. Sci. USA* 96 (18), 10242–10247.
- Lee, K.E., 1985. *Earthworms: Their Ecology and Relation with Soil and Land Use*. Academic Press, London.
- Lenton, T.M., 1998. Gaia and natural selection. *Nature* 394, 439–447.
- Lenton, T.M., 2004. Clarifying Gaia: regulation with or without natural selection. In: Schneider, S.H., Miller, J.R., Crist, E., Boston, P.J. (Eds.), *Scientists Debate Gaia: The Next Century*. London, MIT Press, pp. 15–25.
- Lenton, T.M., Lovelock, J.E., 2000. Daisyworld is Darwinian: constraints on adaptation are important for planetary self-regulation. *J. Theor. Biol.* 206, 109–114.
- Lenton, T.M., Lovelock, J.E., 2001. Daisyworld revisited: quantifying biological effects on planetary self regulation. *Tellus* 53B, 288–305.
- Lenton, T.M., Wilkinson, D.M., 2003. Developing the Gaia theory. *Climatic Change* 58 (1–2), 1–12.
- Lenton, T.M., Caldeira, K.G., Szathmary, E., 2004. What does history teach us about the major transitions and the role of disturbances in the evolution of life and of the Earth system? *Earth System Analysis for Sustainability, Dahlem Workshop Report* 91.H.-J.
- Lovelock, J.E., 1988. *The Ages of Gaia—A Biography of Our Living Earth*. W.W. Norton and Co., New York.
- Lovelock, J.E., 1992. A Numerical model for biodiversity. *Philos. Trans. R. Soc.* 338, 383–391.
- Odling-Smee, J., Laland, K., Feldman, M.W., 2003. *Niche construction: the neglected process in evolution*. Monographs in Population Biology, vol. 37, Princeton University Press.
- Robertson, D., Robinson, J., 1998. Darwinian Daisyworld. *J. Theor. Biol.* 195, 129–134.
- Saunders, P.T., 1994. Evolution without natural selection: further implications of the Daisyworld parable. *J. Theor. Biol.* 166, 365–373.
- Saunders, P.T., Koeslag, J., Wessels, J., 1998. Integral rein control in physiology. *J. Theor. Biol.* 194, 164–173.
- Schwartzman, D.W., Volk, T., 1989. Biotic enhancement of weathering and the habitability of Earth. *Nature* 340, 457–460.
- Silver, M., Di Paolo, E.A., 2006. Spatial effects favour the evolution of niche construction. *Theor. Popul. Biol.* 70 (4), 387–400.
- Staley, M., 2002. Darwinian selection leads to Gaia. *J. Theor. Biol.* 218, 35–46.
- Volk, T., 1998. *Gaia's Body—Toward a Physiology of the Earth*. Copernicus, New York.
- Volk, T., 2004. Gaia is life in a wasteworld of byproducts. In: Schneider, S.H., Miller, J.R., Crist, E., Boston, P.J. (Eds.), *Scientists Debate Gaia: The Next Century*. MIT Press, London, pp. 26–36.
- Watson, A.J., Lovelock, J.E., 1983. Biological homeostasis of the global environment: the parable of Daisyworld. *Tellus* 35B, 284–289.
- Wilkinson, D.M., 1999. Is Gaia really conventional ecology? *Oikos* 84, 533–536.
- Williams, H., Lenton, T.M., 2007. The flask model: emergence of nutrient-recycling microbial ecosystems and their disruption by environment-altering 'rebel' organisms. *Oikos* 116, 1087–1105.
- Williams, H., Noble, J., 2005. Evolution and the regulation of environmental variables. In: Capcarrere, M., Freitas, A., Bentley, J., Johnson, C., Timmis, J. (Eds.), *Proceedings of VIIIth European Conference on Artificial Life, ECAL 2005*, Springer, pp. 332–342.
- Wood, A.J., Ackland, G.J., Lenton, T.M., 2006. Mutation of albedo and growth response produces oscillations in a spatial Daisyworld. *J. Theor. Biol.* 242 (1), 188–198.
- Wood, A.J., Ackland, G.J., Dyke, J.G., Williams, H.T.P., Lenton, T.M., 2008. Daisyworld: A review, *Rev. Geophys.*, 46, RG1001, doi:10.1029/2006RG000217.