

5

Gaia and Evolution

Timothy M. Lenton and Hywel T. P. Williams

We present in this chapter a search for Gaia in computer-generated model worlds. The computer may seem like an odd place to be looking for a planetary-sized phenomenon when we could be examining the real world. However, with a sample size of only one Earth the inferences that can be drawn about the likelihood of certain features are necessarily limited. In particular, our existence as observers who can look back on and marvel at Earth history is only consistent with a history in which atmospheric oxygen built up to the levels necessary to support animals with large brains (Watson 2004) and the climate became (or remained) relatively cool (Schwartzman 1999). One can imagine myriad other scenarios for Earth-like planets in which life never reaches the stage of conscious observers, but by definition no observer is there to see such histories. In the next few decades we may be fortunate to learn about the atmospheric composition of planets orbiting other stars, and from that information we may learn something about the presence or absence of life on them (Lovelock 1965). This could increase the sample size of inhabited planets above one. In the meantime, by creating many virtual worlds in the computer, we can begin to examine whether features we see on Earth, such as abundant recycling and environmental regulation, are likely or unlikely phenomena once life has emerged on a planet.

Simulating Gaia

The search for “Gaia in the machine” (Downing 2004) began with Daisyworld (Watson and Lovelock 1983; Lovelock 1983a), and there have been many variants of it since (Wood, Ackland, et al. 2007). Perhaps because of Daisyworld’s elegant simplicity and great adaptability, relatively few alternative model worlds have been developed.

Notable exceptions are the Guild (Downing and Zvirinsky 1999) and Metamic (Downing 2002) models, although these share some key assumptions with Daisyworld. There are also stripped-down models based on Daisyworld that remove many of its key assumptions (Staley 2002; McDonald-Gibson, Dyke, et al. 2008). The danger with modeling is that it is always possible to make a model to illustrate a particular point. However, it is now possible to make stochastic, evolutionary models with very many degrees of freedom and a huge range of possible outcomes, many of which cannot be anticipated by the modeler. This is the approach we take here, of “testing Gaia theory with artificial life” (Lenton 1999). In this respect we follow the pioneering work of Keith Downing who first applied “artificial life” techniques to tackle Gaia questions (Downing and Zvirinsky 1999; Downing 2002, 2004).

Gaia theory posits that a Gaia system will tend to self-regulate in a habitable state, one in which life can survive. Here we follow our previous definitions of a Gaia system, using the terminology of regulation and self-regulation (Lenton 2004). A “Gaia system” is a type of planetary-scale, open thermodynamic system, with abundant life supported by a flux of free energy from a nearby star. *Regulation* describes the return of a variable to a stable state after a perturbation. *Self-regulation* describes a system automatically bringing itself back to a stable state, rather than an external agent imposing regulation or any conscious purpose (teleology) within the system bringing about regulation. For internal consistency of the definitions, and to distinguish a Gaia system from a planet in which, for example, a few extremophiles survive below the surface, we narrow the definition of “habitable state” to one which supports abundant life. Here we also consider nutrient recycling, which can be defined as occurring when the flux of a given nutrient through primary producers exceeds the input flux of that nutrient into a system (Volk 1998).

Why Evolutionary Biologists Don’t (or Didn’t) Like Gaia

We begin by reviewing the evolutionary critiques of Gaia. At first glance one might have expected biologists to greet with enthusiasm the idea that life plays a major role in the regulation of the planet. After all, it gives extra prominence to their subject area. Of course, eminent microbiologist Lynn Margulis championed Gaia (Margulis and Lovelock 1974; Lovelock and Margulis 1974a; Lovelock and Margulis 1974b), and she

was not (completely) alone in confronting the many neo-Darwinian evolutionary theorists who reacted with vehement opposition when Gaia was first proposed (Doolittle 1981; Dawkins 1983). While generally recognizing that the Earth does display some remarkable stabilizing properties, the critics contended that there was no mechanism by which “selfish” genes and organisms could come to regulate the planet. In other words, Gaia may work in practice but it will never work in theory!

The first protest was that the notion of “atmospheric homeostasis by and for the biosphere” (Lovelock and Margulis 1974a) implies teleology—some conscious foresight or planning on the part of unconscious organisms. This was convincingly answered by the Daisyworld model (Watson and Lovelock 1983; Lovelock 1983a), which shows that self-regulation can occur without teleology in a feedback system of life coupled to its nonliving environment. This should surprise no one trained in systems theory. In any coupled system with a mixture of positive and negative effects forming multiple feedback loops there is a good chance that the system will settle down in a negative feedback regime. However, most evolutionary biologists (in contrast to their now distant colleagues in physiology) seem blissfully unaware of systems theory. As Jim Lovelock once put it (in conversation with T.M.L.); “I know professors of biology who have trouble with the concept of self-regulation, but they have no difficulty walking!”

The second protest was that while the self-regulating properties of organisms have been refined by natural selection, the Earth exists in a population of one, and therefore any self-regulation it displays cannot have been shaped by natural selection at the global scale. Given the obvious power of natural selection to engineer well-adapted individuals, this may imply that self-regulation in organisms will be more finely honed than any seen at the planetary scale. However, it does not deny the existence of regulation at the planetary scale. Even in organisms, natural selection cannot *create* self-regulation; it can only favor those individuals that happen to display self-regulation in the sense that they leave more descendants. Evolution by natural selection really comprises three parts: the tendency toward exponential growth (creating competition and selection pressure), heritable variation based on (near) faithful replication, and some source of innovations (new traits) that give differential survival rates. The “innovation” of self-regulation need not stem from a point mutation at some genetic locus; it may simply be the

tendency of a sufficiently complex feedback system to settle in a negative feedback regime, a mechanism that is just as applicable to the planet as it is to organisms.

The third protest was that any system in which certain “altruistic” life forms expended energy contributing to making a better global environment would be vulnerable to “cheats”—organisms that enjoyed the benefits of a better global environment but did not contribute to it. Cheats would save energy and thus outcompete altruists, ultimately destroying the regulatory system. The limitation of this argument is that it is predicated on the ideas that organisms can “choose” whether or not to alter the environment and that it will cost energy to do so. In fact as anyone aware of thermodynamics should know, altering the environment is an unavoidable consequence of being alive and can be part of a process that transfers energy to the organism (Schrödinger 1944). Living organisms are highly ordered (low-entropy) structures, and to maintain an ordered state, they must take in matter and free energy (often combined together as “food”), transform the matter, and excrete waste products that are of a lower free energy (higher entropy) state, with some energy always degraded as heat. Plants and other photosynthesizers are an exception in that they can use the free energy in sunlight to turn lower energy compounds from the environment into higher energy compounds, but even they must then break down these low-entropy compounds to fuel their own metabolism.

An evolutionary biologist would rightly counter that in addition to the inevitable change of the environment that comes about with metabolism, most forms of Gaia (early hypothesis or later theory) postulate that life forms have traits that have been selected for their environment-altering properties. Daisyworld clearly invokes such traits—the blackness or whiteness of daisies in Daisyworld is (presumably) not an inevitable consequence of their being alive. These traits alter the local temperature of each daisy and its progeny, and they also alter the global temperature in the same direction. This means that traits that are good or bad at the level of individual selection are correspondingly good or bad for global temperature regulation (relative to the shared optimum of all the daisies). As many evolutionists have pointed out, this makes the model a special case, and one that is “designed” to give regulation. Many variants of the model have been created and they generally show regulation (Wood, Ackland, et al. 2007). The main exception is when the optimum growth temperature of the daisies is allowed to adapt unbounded, but that is

an unrealistic scenario given the thermodynamic constraints on real biochemistry (Lenton and Lovelock 2000).

Interestingly, when the system is rewired so that black daisies produce white clouds that cool the global environment, regulation still occurs (Watson and Lovelock 1983). This is far from being an altruistic world—the white daisies are outcompeted and barely get a look in. However, evolutionary biologists have persisted in seeing Gaia as involving some form of altruism, and since William Hamilton's seminal work in the 1960s it has been clear that the conditions under which altruism can flourish are rather restrictive (Hamilton 1964, 1972). Furthermore, when Gaia hit the popular consciousness in the late 1970s, neo-Darwinists were struggling to rid their subject of arguments predicated on group selection. Gaia appeared to be perhaps the most extravagant example of altruism demanding higher level selection, consequently it was lumped in with the worst examples of arguments “for the good of the species,” and summarily dismissed. In the last few years there has been a resurgence of interest in the evolution of cooperation (Nowak 2006) and multilevel selection (Sober and Wilson 1997; Goodnight 2005; Bijma, Muir, et al. 2007). While by no means all of the evolutionary biology community approves of explanations involving selection above the level of the individual (or even above the gene), they are at least up for open discussion again. Thus, in some sense, Lovelock and other early proponents of Gaia were unlucky in their timing.

The final reason why evolutionary biologists and many other scientists don't like Gaia is the mythological name. Opinion varies considerably from scientist to scientist—for example, John Maynard Smith described it (in conversation with T.M.L.) as “an awful name for a theory,” but Bill Hamilton had no objection and was happy to write a paper with “Gaia” in the title (Hamilton and Lenton 1998), so long as it was recognized that a Gaia system is a different class of system than an organism or a superorganism such as a termite mound. Unfortunately, Lovelock's likening of Gaia first to an organism and then to a superorganism agitated evolutionary theorists like Hamilton and Maynard Smith, because they wished to reserve the terms “organism” and “superorganism” for systems that can show adaptations due to natural selection. Kin selection can help explain the system level properties of the wasp nest and the bee colony, and it is for such systems that evolutionary biologists want to reserve the term “superorganism” (Hamilton 1964; Maynard Smith 1964).

Current Status of the Debate

The crux of the present Gaia debate is: Do we need to invoke individual-level traits that have been selected for their environment-improving properties in order to account for observed Gaian properties of the global system? (And if we do, how can we avoid the problem of “cheats”?) Or, can we construct a reasonable theory of regulation based entirely on environment-altering properties that are simply by-products of metabolic traits selected for other reasons? The concept of by-products was independently introduced to the Gaia debate in the late 1990s by three different authors—Lenton (1998), Volk (1998), and Wilkinson (1999); and it was Volk (1998) who promoted the term “by-product.” If regulation arises in a system built entirely on by-products, this makes Gaia theory much less vulnerable to criticisms from evolutionary theory. However, some authors are skeptical that regulation via by-products can occur, emphasizing recycling instead as the key Gaian property (Volk 1998).

To help understand the mechanisms at play in global regulation, one of us (T.M.L.) introduced the distinction between “feedbacks on growth” and “feedbacks on selection” (Lenton 1998). Feedback on growth occurs when a trait alters the environment in a way that affects the growth of its carriers and noncarriers in equal measure such that there is no change in the forces of selection. Feedback on selection occurs when a trait alters the environment in a way that affects the growth of its carriers and noncarriers differentially, and thus influences its own selection. In either case the responsible trait may initially arise as a selectively neutral by-product. In the case where the trait generates only feedback on growth it will remain a selectively neutral by-product, since it can offer no individual-level selective advantage. However, if the trait generates feedback on selection, it may become adaptive or maladaptive depending on its environmental effects, and its neutral by-product status is lost.

In ecology the related concepts of “extended phenotype” (Dawkins 1983), “niche construction” (Odling-Smee, Laland et al. 2003), and “ecosystem engineering” (Jones, Lawton, et al. 1994) involve the idea that genes or organisms can shape their abiotic environment in a way that alters their selective environment. We think feedback on selection may be most relevant in systems with environmental heterogeneity at such intermediate scales. However, we begin our modeling by removing the possibility of feedback on selection.

Our Search for Gaia

Our approach has been to build a system where environmental alteration is entirely a no-cost by-product of metabolism, and to see what conditions lead to the emergence of nutrient recycling and regulation. We base our model around bacteria, since they are ubiquitous, have a 3.8 billion year pedigree, run the global biogeochemical cycles, and are highly adaptable. For reasons of tractability (of analysis as well as computation) we look at a model microcosm, rather than the global macrocosm, but the microcosm approach has other advantages. We believe that similar principles must be at play in systems of all scales so that we can learn something useful about the macrocosm from the microcosm. Microcosm studies also offer the possibility of empirical testing in the laboratory, something that would clearly be impossible with a global model. We incorporate evolution by modeling individuals that each have a genetic code, and allowing selection pressures on these individuals to emerge from the dynamics of the microbial ecosystem.

The flask model (figure 5.1) simulates an evolving population of microbes suspended in a flask of liquid, and hence the name (Williams and Lenton 2007a, b). There is a prescribed supply flux of different nutrients into the flask and corresponding removal fluxes proportional to the concentration of each nutrient in the flask. There are also non-nutrient “abiotic” environmental variables. The flask is seeded with a clonal population of model microbes. Each microbe has a genetic code that prescribes its uptake and release patterns for nutrients, its preference for the abiotic environment, and its metabolic by-product effect on abiotic environmental variables. We place only one genetic constraint on metabolism: an organism is not allowed to consume what it excretes. However, an important constraint on metabolism is built into the model in the form of a peaked metabolic rate function that declines smoothly to zero as the state of the environment moves away from optimal growth conditions. Reproduction is asexual, and at each reproduction event there is the possibility of random mutation. We remove the possibility that individuals can differentially benefit from altering their environment by assuming that the liquid environment inside each flask is well mixed such that any environmental change is experienced equally by all individuals.




Figure 5.1
Schematic of the flask model.

The Emergence and Disruption of Nutrient Recycling

The emergence of nutrient recycling loops is a robust result in our model system (Williams and Lenton 2007a). We measure recycling ratios (Volk 1998; Downing and Zvirinsky 1999) as the biological uptake flux of a given nutrient at each time step divided by the input flux of that nutrient to the flask. A value of 1 indicates efficient uptake of that nutrient. A value greater than 1 indicates that recycling is occurring. Initially we ran the model without any constraints on growth from the abiotic environment (and hence no environmental feedback) in order to concentrate on the emergence of nutrient recycling.

Selection acting on nutrient consumption traits causes the community to be dominated by specialist organisms adapted to take up a single nutrient (Williams and Lenton 2007a). This occurs because organism growth depends on the total quantity of nutrients consumed. Different nutrients are always taken up in fixed proportions set by the genetic code

of each individual such that overall consumption is limited by scarcity of any required nutrient. Specializing on a single nutrient minimizes the risk of limitation. Nutrient release patterns do not directly affect the fecundity of individuals; therefore there is no selection of release traits. Organisms typically excrete a mixture of nutrients (excluding the one they consume).

In a typical run with four nutrients (figure 5.2), after initialization one nutrient is efficiently taken up and the population is limited by its supply. After around 8,000 time steps, adaptation of consumption traits allows a second nutrient to be efficiently utilized, some nutrient recycling of the two fully utilized nutrients begins, and the population increases. When further adaptation allows a third nutrient to be efficiently utilized (after around 18,000 time steps), the population size and recycling ratios rise again, but it is the efficient uptake of the fourth nutrient (after around 30,000 time steps) that really boosts the population size and recycling ratios. The population then reaches a carrying capacity that is set by the prescribed inputs of nutrients and the assumption of a fixed fraction of energy being lost as heat in metabolism.

When we introduce constraints on growth from the abiotic environment this leads to selection on the environmental preferences of the organisms (Williams and Lenton 2007a). Once again single-nutrient consumers dominate and efficient uptake and recycling of all nutrients becomes established. Typically, after some initial meandering, the environment settles in a fairly stable state and the community converges on shared preferences for that environment. There is no selection of abiotic effects on the environment because these cannot give differential benefit to individuals (because of the well-mixed shared environment). Consequently this part of the gene pool shows high diversity and genetic drift. Genetic drift can cause the environment to shift to a different state and the population's environmental preferences adapt in response.

As we tighten the environmental constraints on growth a new phenomenon emerges—the system becomes vulnerable to population crashes (figure 5.3). Although “cheating” (in the sense meant by evolutionary biologists) is not possible in a system built on by-products, “rebel” species can appear that disrupt the system by rapidly shifting the environment into an inhospitable state as they exploit previously unused nutrient stocks (Williams and Lenton 2007a). Vulnerability to such changes is worsened by genetic convergence of the population around a shared preferred environmental state. This convergence can cause a population

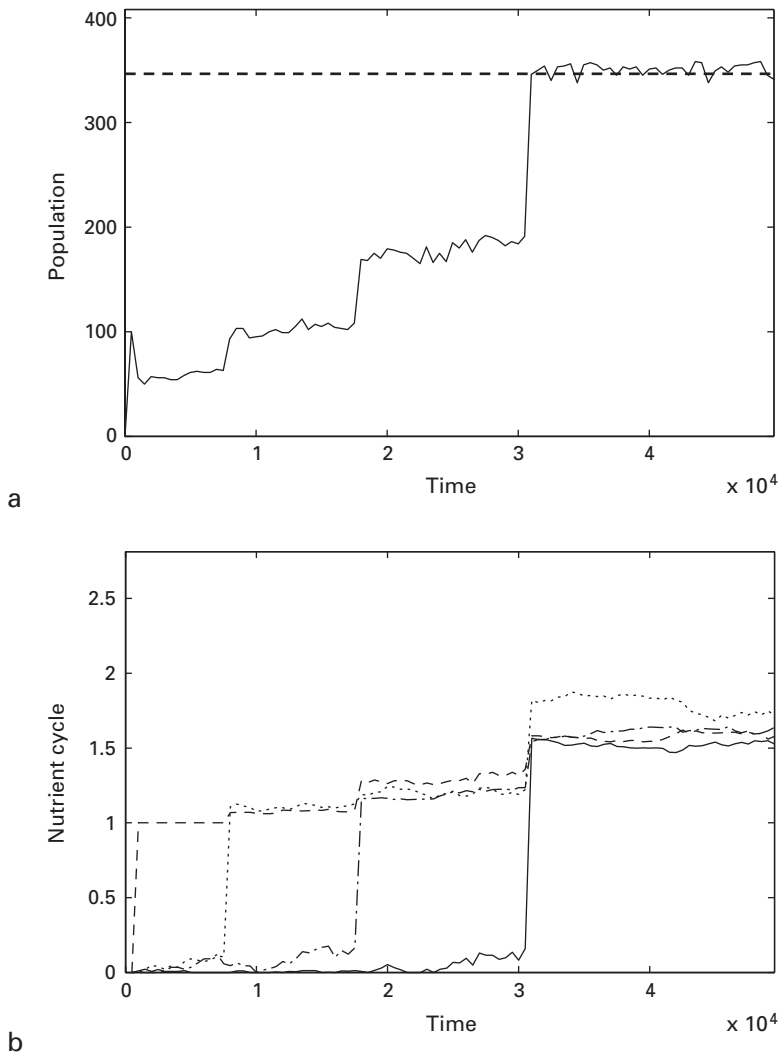


Figure 5.2

Establishment of nutrient recycling in a typical run of the flask model with no constraints from the abiotic environment: (a) Population size (solid line) and analytically derived carrying capacity (dashed line), (b) nutrient recycling ratios for the four nutrients in the system (each with a unique line type). Time is measured in 10^4 time-steps (the plots show 50,000 time-steps). Image adapted from Williams and Lenton (2007a).

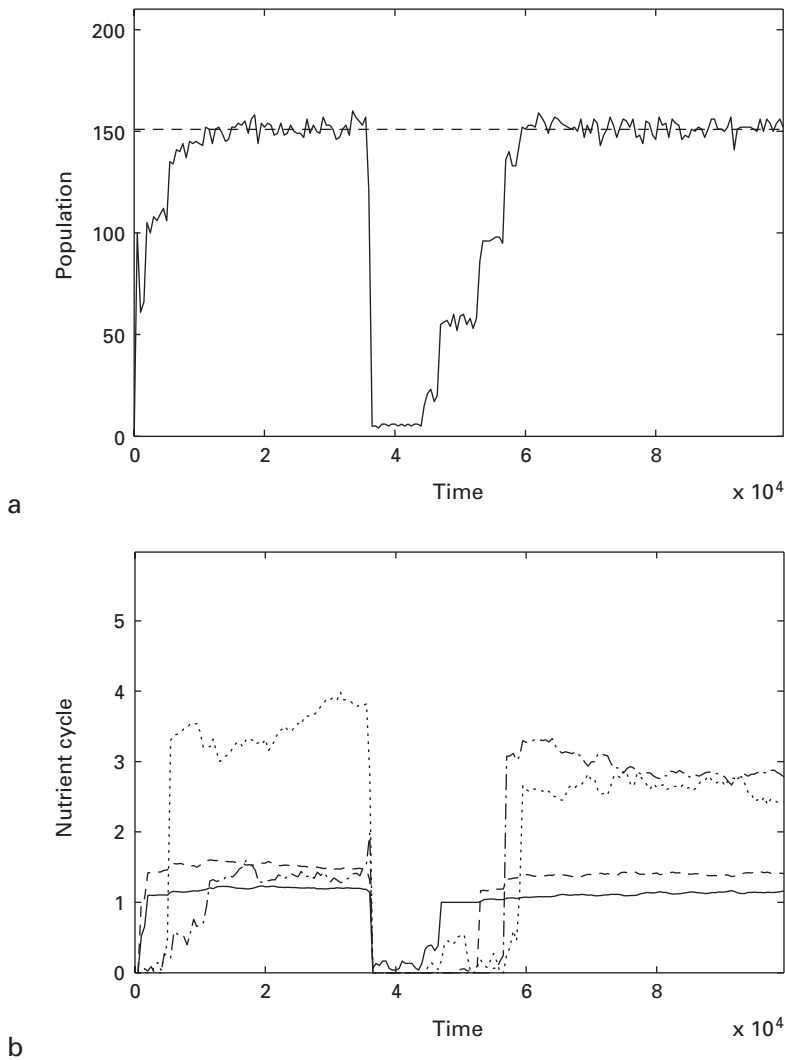


Figure 5.3

Establishment, collapse, and subsequent recovery of nutrient recycling in a run of the flask model with relatively tight environmental constraints on growth: (a) Population size (solid line) and analytically derived carrying capacity (dashed line), (b) nutrient recycling ratios for the four nutrients in the system (each with a unique line type). Time is measured in 10^4 time-steps, (the plots show 100,000 time-steps). After around 35,000 time-steps there is a crash in population size and nutrient recycling caused by a “rebel” organism that utilizes an underused nutrient whilst shifting the environment away from the state to which the population is adapted. Image adapted from Williams and Lenton (2007a).

crash when the environment moves too far from the condition to which the population is adapted. Sometimes the population gradually adapts to the new environmental state and efficient nutrient consumption and recycling re-emerge, but in the worst case, extinction can occur. In this latter case members of the “rebel” species destroy the ecosystem and themselves with it—like an overvirulent parasite.

We ran large ensembles of runs for many versions of the model to quantify the robustness of the results (Williams and Lenton 2007a). The emergence of recycling is a robust result regardless of the tightness of constraints from the abiotic environment. The average recycling ratio typically asymptotes after around 50,000 time steps, at a value that depends on the efficiency of conversion of nutrients into biomass during metabolism. For example, for a nutrient conversion efficiency of 60 percent, recycling ratio asymptotes at around 2, meaning that approximately the same amount of consumed material is from nutrient recycling as from external supply. Extinction rates increase with tightened constraints from the abiotic environment (figure 5.4). These rates count both

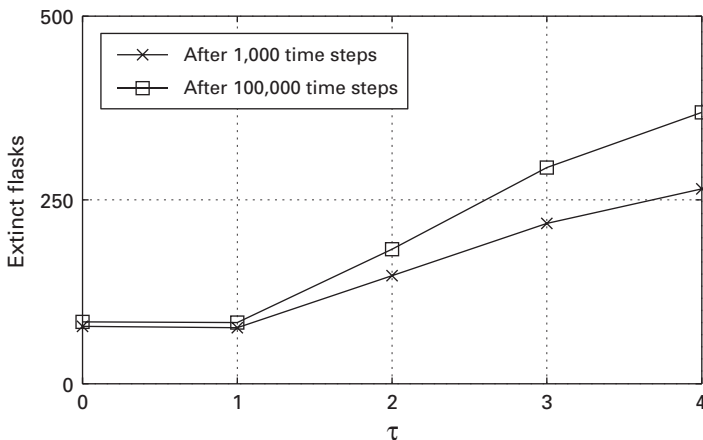


Figure 5.4

Number of extinct flasks in an ensemble of 500 runs as a function of the tightness of environmental constraints on growth (determined by the parameter τ in the model). The line with crossed markers shows initial extinctions due to communities being mismatched to their environment. The line with square markers shows total extinctions at the end of a long run. The difference between the lines indicates the internally generated extinctions due to “rebel” organisms during the model runs. The number of initial extinctions and internally generated extinctions both increase with tightening environmental constraints on growth.

extinctions at initialization because the seed community is unsuited to its environment and endogenously created extinction events caused by rebel species. However, even for very tight constraints from the abiotic environment, a significant fraction of systems survive.

In a single flask the abiotic environment can exhibit long intervals of quasi-stability with relatively small variation, punctuated by brief episodes of large shifts in state. However, this is not true regulation because the systems do not recover from perturbations. The dominant behavior in the single-flask system is robust nutrient cycling with stochastic variation in environmental state caused by genetic drift in the environment-altering traits of the population, sometimes with occasional discontinuities caused by population crashes when abiotic constraints on growth are tight.

The Emergence of Environmental Regulation

To continue our search for environmental regulation in our virtual worlds, we constructed a spatial version of the flask model with multiple interconnected flask ecosystems (Williams and Lenton 2008). Typically we have ten flasks arranged in a ring with mixing between nearest neighbors. Mixing occurs at each time step by transferring a fixed volume of liquid between adjacent flasks. (Imagine simultaneously dipping and swapping cups of liquid from neighbor to neighbor.) This simple approximation to diffusive mixing transfers nutrients, abiotic factors, and organisms between the flasks. We retain perfect mixing within each flask but vary the degree of mixing between the flasks. Imperfect mixing between flasks introduces environmental heterogeneity into the global system: although the mixing process would eventually homogenize the global environment, it does not happen on a fast enough timescale to overcome the continual, differentiating metabolic activities of the local microbe populations.

We have explored this system with and without adaptation of the environmental preferences of the organisms. The same qualitative results emerge but the interpretation is much easier in the case where we switch off adaptation of environmental preferences and give all organisms the same fixed environmental preference. We measure the “environmental error” as the difference between the average environmental state and the shared preference of the organisms. Thus with fixed preferences, any reduction in environmental error must be due to the organisms collectively shifting the environment toward their preference. Local

endogenous extinctions can still occur within individual flasks, as they did in the single-flask system. However, the rebel species responsible rarely succeed in spreading to destroy the global system. It is more usual to see instead the denuded flask being rapidly recolonized from its neighbors—an example of metapopulation dynamics (Hanski 1998).

More interestingly, for intermediate mixing rates we see a reduction in environmental error over time (figure 5.5), even when the system is perturbed by periodically changing the input fluxes of nutrients and the level of abiotic factors. The environmental error in an individual run is not reduced to zero, but instead the system undertakes behavior




Figure 5.5

Case study of the emergence of environmental regulation in a spatial system of 10 flask ecosystems connected in a ring. The solid line with error bars denotes the actual state of the abiotic environment (mean value across all 10 flasks). The dashed line indicates the calculated abiotic environmental state in the absence of life. The dash-dot line indicates the universally shared microbial preference for the abiotic environment. The dotted lines indicate the bounds of the habitable range, found where metabolism exactly balances the maintenance costs of living. The system is subject to random perturbations to the external forcing of the abiotic environment every 5,000 time-steps. The mismatch between the actual state of the environment and the preferred state generally reduces over time and the system counteracts perturbations.

approximating to bounded stochastic variation between the upper and lower limits of habitability. Here the bounds of habitability correspond to the environmental states at which metabolism brings in just enough energy to meet the essential maintenance costs of being alive. Across an ensemble of runs (figure 5.6), the mean error asymptotes to a value that represents the average distance from optimality of a system that is undertaking a random walk between the habitability bounds.

Measurements of the variation of several system metrics (population size, growth rate, nutrient availability) against environmental error suggest that we have a system with two regimes, one in which nutrients limit growth and the environment has no effect, and one in which the environment limits growth and nutrients are abundant. The environment-limited regime exists at and outside the bounds of habitability. When the




Figure 5.6

Reduction in environmental error over time averaged over ensembles of 500 unperturbed runs (solid line) and 500 perturbed runs (dashed line) (where figure 5.5 is an example of a perturbed run). The “error” is simply the magnitude of the difference between the preferred and the actual state of the abiotic environment. The dot-dash line indicates the expected size of the environmental error in the absence of life. The dotted line indicates the boundary of the habitable range, that on average the system would be in an uninhabitable state in the absence of life.

system strays outside these bounds then any ecosystem that collectively improves the environment (drags it back toward the habitable region) will experience positive feedback on growth, whereas any ecosystem that collectively degrades the environment (pushes it further away from the habitable region) will experience negative feedback on growth. However, this does not explain why we typically observe environment-improving local communities dominating the global system.

The observed dominance of environment-improving communities requires a mechanism by which they outcompete environment-degrading communities (figure 5.7). In the spatial flask model, this mechanism is selection at the level of the local ecosystem based on their differential rates of proliferation. Simply put, environment-improving communities become larger because they reduce the limiting environmental constraint on their growth, while environment-degrading communities become

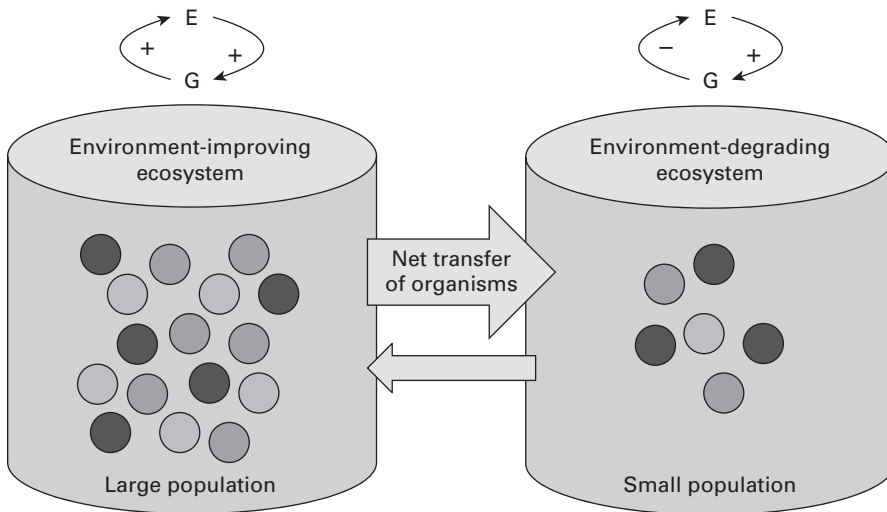


Figure 5.7

Schematic of the mechanism by which environment-improving ecosystems tend to dominate the global system. In the feedback loops at the top, E denotes the environmental state and G denotes growth. An ecosystem that collectively improves its environment generates positive feedback on growth increasing its population. An ecosystem that degrades its environment generates negative feedback on growth restricting its population size. Mixing between flasks occurs by an exchange of equal volumes of fluid and whatever microbes they contain. Thus the environment-improving ecosystem with higher population density tends to spread more rapidly than the environment-degrading ecosystem with lower population density.

smaller. Larger (and thus denser) communities spread more effectively than smaller ones because the fixed-volume transfer mechanism between flasks carries more individuals from a higher density population source than it does from a lower density population. This imbalance allows the members of environment-improving communities to eventually take over the global system, shifting the global environment toward the optimum for growth.

The hard-headed evolutionary biologist might suspect that the phenomena we observe in our model do not involve higher level selection operating on whole ecosystems, and that they may instead be explained by the implicit selection of a single “super-species” within an ecosystem that is alone responsible for improving the environment. Indeed we have seen such super-species in artificial ecosystem selection experiments with the flask model (Williams and Lenton 2007b). However, inspection shows that in the spatial system local communities are always highly diverse, with many species interacting in complex ways subject to individual level selection. There is high local diversity in values for the environment-altering traits, showing that these traits are selectively neutral at the individual level, but the global mean value of environment-altering traits clearly alters to counteract changes in external forcing. This demonstrates that selection pressure on these traits is active at some level, which our analysis shows to be the level of the ecosystem. We do not claim any long-term community-level adaptation (given the ongoing disruption from migration and mutation) but only the presence of selection acting over short timescales to promote the spread of communities that improve their environment over those that degrade their environment.

Where Next?

Our search for Gaia in the computer appears to have succeeded. The multiple-ecosystem flask model self-regulates (in the sense that it maintains a habitable environment and counteracts perturbations), despite being built on by-products. Ensembles of runs with the spatial model show a progressive improvement in the environment over time. These results are especially interesting given arguments by others questioning whether a global system built on by-products would tend to regulate (Volk 1998). The mechanism of regulation involves multiple levels of selection without needing selection at the level of the planet, and we believe that it represents a novel mechanism for generating environmental regulation. Although reproduction of an ecosystem is clearly less

faithful in its replication than reproduction of an organism, there is still enough short-term heritability of collective environment-improving properties for communities with them to spread across the global system. Interestingly our mechanism involves shifting the localization of the environment to the ecosystem level so that our flask ecosystems become somewhat akin to the daisies of Daisyworld.

We have run the model with evolvable preferences and obtain qualitatively similar results. It is harder to disentangle what is causing the behavior when preferences evolve. This is because moving preferences toward the environment has an environmental error-reducing effect equivalent to that of moving the environment toward preferences. We have done extensive parameter sensitivity studies of the model system in both scenarios, and the results are robust. We have also reflected on how evolutionary biologists might critique the model. Clearly, we have built into the spatial model the necessary level of population structure for multilevel selection to operate. Environmental heterogeneity at some scale and corresponding selection of communities based upon it is critical to getting environmental regulation. However, such structure and environmental heterogeneity exist in the real world.

Having homogenized the local environment in the model, and banished environment-altering traits that cost their carriers, one next step would be to relax these assumptions. A more general model would allow the possibility of individual-level environmental alteration and selection of the traits responsible. Our contention is that there are cases in the real world where costly environment-altering traits are selected because they beneficially alter the immediate environment of the organism sufficiently to outweigh the cost and bring a net fitness increase. This scheme may be extended to improving the environment of one's offspring or other relatives, provided that the cost of such altruism is balanced by the benefit accruing to kin (Hamilton 1964). Niche construction (Odling-Smee, Laland, et al. 2003) and its close relation, ecosystem engineering (Jones, Lawton, et al. 1994), describe cases where the environment-altering activities of an organism alter the selection pressures faced by the individual and its descendants. Examples where costly environmental alterations offer a selective benefit include the beaver's dam, ant colonies, termite mounds, and many more (Odling-Smee, Laland, et al. 2003). There are also a host of phenomena in terrestrial ecosystems, such as sphagnum moss forming a peat bog that excludes trees, where it is unclear which mechanism to invoke to explain what is going on.

Hamilton argued (personal communication to T.M.L.) that multilevel selection might play a role in shaping such systems. By allowing for a wider range of possibilities in the model, we could address what is the most likely explanation for specific scenarios in the real world.

Wider Implications

We have taken the constructive critics of Gaian regulation seriously and built a model with environmental alteration based only on by-products of metabolism. Recycling emerges as a robust Gaian property and in our spatial system environmental self-regulation also occurs. This suggests that what critics argue is the most “acceptable” version of Gaia theory may need revising. Even if “Gaia is life in a wasteworld of by-products” (Volk 2004), such a Gaia can still self-regulate, keeping the environment within habitable bounds and counteracting perturbations. A necessary condition for regulation is heterogeneity in the environment at some scale (in our model, between local flask ecosystems). This means we cannot apply the argument to truly well-mixed global variables such as the dominant gases in the atmosphere, but in principle, it could apply to any biologically-influenced non-nutrient variable that exhibits spatial gradients between different parts of the planet, such as aspects of the climate (e.g., temperature).

We also suggest that one should not be too dogmatic in focusing only on the version of the Gaia theory that is acceptable to critics—we strongly suspect that there are more than just by-products at play in shaping the real Gaia system. It is conceivable that an environment-altering trait may start life as a by-product but then be selected for its environmental effects, thus becoming adaptive. This was the argument put forward by Hamilton and one of us (T.M.L.) when thinking about the benefits of dimethyl sulphide emission and biogenic ice nucleation to aerial dispersal of spora (Hamilton and Lenton 1998).

In our model worlds, seemingly “cooperative” enterprises such as recycling and regulation robustly appear. In the spatial system, when the environment is limiting to growth, ecosystems or communities that “foul their nest” lose out to those that improve their local conditions. This has implications beyond our model, such as for life on the early Earth, for modern ecosystems, and for the human situation. There is no “tragedy of the commons” (Hardin 1968) in our model worlds. All individuals change their environment, but when the environment is

limiting, environment-improving ecosystems or communities come to dominate.

It is pertinent to ask whether any parallels can be drawn between our model mechanism and the current predicament of human communities causing and experiencing climate change. Up until now, we humans have been a “rebel” species, altering the global environment as a by-product of more locally selected activities; for example, carbon dioxide emissions are a by-product of fossil fuel burning to produce energy. The “community-selection” mechanism cannot help solve the problem of rising carbon dioxide levels because carbon dioxide is a globally well-mixed variable. In this case the “tragedy of the commons” applies: communities that lower their carbon dioxide emissions will not see a differential climate benefit to those that increase theirs. However, if communities use localized mechanisms to lower their temperature, then there may be some scope for them to feel a differential benefit. Already there is inadvertent aerosol cooling due to fossil fuel (especially dirty coal) burning, which is masking greenhouse warming across some regions. If aerosol cooling were taken away, these regions would experience the greenhouse warming unmasked. In California, measures are being implemented to make building roofs and vehicles more reflective in order to counteract the combined global warming and urban heat island effects. More radically the prospect of climate engineering has been proposed, in particular, injecting aerosol into the stratosphere to cool the surface. If nations, for example, chose to deploy such strategies for their own benefit, this could introduce interesting dynamics in the global system. Whether such strategies are “selected” depends on their cost relative to the savings from avoided climate change damages. However, it is now clear that climate change damages could be expensive enough to restrict the growth of economies (Stern 2006), raising the possibility that a form of economic “community selection” could occur. The communities selected would be those that deploy strategies to keep themselves cool that cost less than the climate damages they manage to avoid. The community could be at a range of scales, from villages to continents, depending on the strategy used.

Conclusion

Our search for Gaia in the computer has proved illuminating. In our virtual worlds, recycling and self-regulation of the environment within habitable bounds robustly emerge. If we draw an analogy between

our model and microbial life on the early Earth, then we have some confidence that once life got started it would have soon solved the problem of nutrient recycling. Once there was a population of bacterial ecosystems, perhaps in ponds in different meteorite craters, or in different gyres of the ocean, or around different hydrothermal vents, environmental self-regulation could have emerged. Self-regulation would have worked best for somewhat heterogeneous environmental variables. We are still a long way from addressing the question (raised at the start of this chapter) of how probable a regulating and recycling Gaia system like the one we inhabit today might be, given the existence of life on a planet and 4 billion years of evolution. It is also important to remember the limitations of computer modeling and that artificial life *in silico* is fundamentally different from biological life. However, with these caveats in place, we nonetheless have increased confidence that a self-regulating “microbial Gaia” is a probable outcome once life gets started on a planet.

References

- Bijma, P., W. M. Muir, and J. A. M. Van Arendonk. 2007. Multilevel selection. 1: Quantitative genetics of inheritance and response to selection. *Genetics* 175: 277–88.
- Dawkins, R. 1983. *The Extended Phenotype*. Oxford: Oxford University Press.
- Doolittle, W. F. 1981. Is Nature really motherly? *CoEvolution Quarterly* (spring): 58–63.
- Downing, K., and P. Zvirinsky. 1999. The simulated evolution of biochemical guilds: Reconciling Gaia theory and natural selection. *Artificial Life* 5 (4): 291–318.
- Downing, K. L. 2002. The simulated emergence of distributed environmental control in evolving microcosms. *Artificial Life* 8 (2): 123–53.
- Downing, K. L. 2004. Gaia in the machine: The artificial life approach. In S. H. Schneider, J. R. Miller, E. Crist, and P. J. Boston, eds., *Scientists Debate Gaia: The Next Century*. Cambridge: MIT Press, pp. 267–80.
- Goodnight, C. J. 2005. Multilevel selection: The evolution of cooperation in non-kin groups. *Population Ecology* 47: 3–12.
- Hamilton, W. D. 1964. The genetical evolution of social behaviour. *Journal of Theoretical Biology* 7: 17–52.
- Hamilton, W. D. 1972. Altruism and related phenomena, mainly in social insects. *Annual Review of Ecology and Systematics* 3: 193–232.

- Hamilton, W. D., and T. M. Lenton. 1998. Spora and Gaia: How microbes fly with their clouds. *Ethology Ecology and Evolution* 10: 1–16.
- Hanski, I. 1998. Metapopulation dynamics. *Nature* 396: 41–49.
- Hardin, G. 1968. The tragedy of the commons. *Science* 162: 1243–48.
- Jones, C. G., J. H. Lawton, and M. Shachak. 1994. Organisms as ecosystem engineers. *Oikos* 69: 373–84.
- Lenton, T. M. 1999. Testing Gaia theory with artificial life. In D. Floreano, J.-D. Nicoud, and F. Mondada, eds., *Proceedings of the 5th European Conference on Artificial Life*. Berlin: Springer, pp. 9–10.
- Lenton, T. M. 1998. Gaia and natural selection. *Nature* 394: 439–47.
- Lenton, T. M. 2004. Clarifying Gaia: Regulation with or without natural selection. In S. H. Schneider, J. R. Miller, E. Crist, and P. J. Boston, eds., *Scientists Debate Gaia: The Next Century*. Cambridge: MIT Press, pp. 15–25.
- Lenton, T. M., and J. E. Lovelock. 2000. Daisyworld is Darwinian: Constraints on adaptation are important for planetary self-regulation. *Journal of Theoretical Biology* 206 (1): 109–14.
- Lovelock, J. E. 1965. A physical basis for life detection experiments. *Nature* 207 (4997): 568–70.
- Lovelock, J. E. 1983a. Gaia as seen through the atmosphere. In P. Westbroek and E. W. de Jong, eds., *Biom mineralization and Biological Metal Accumulation*. Dordrecht: Reidel, pp. 15–25.
- Lovelock, J. E., and L. M. Margulis. 1974a. Atmospheric homeostasis by and for the biosphere: the Gaia hypothesis. *Tellus* 26: 2–10.
- Lovelock, J. E., and L. M. Margulis. 1974b. Homeostatic tendencies of the Earth's atmosphere. *Origins of Life* 5: 93–103.
- Margulis, L., and J. E. Lovelock. 1974. Biological modulation of the Earth's atmosphere. *Icarus* 21: 471–89.
- Maynard Smith, J. 1964. Group selection and kin selection. *Nature* 201: 1145–47.
- McDonald-Gibson, J., J. G. Dyke, E. A. Di Paolo, and I. R. Harvey. 2008. Environmental regulation can arise under minimal assumptions. *Journal of Theoretical Biology* 251: 653–66.
- Nowak, M. A. 2006. Five rules for the evolution of cooperation. *Science* 314: 1560–63.
- Odling-Smee, F. J., K. N. Laland, and M. W. Feldman. 2003. *Niche Construction: The Neglected Process in Evolution*. Princeton: Princeton University Press.
- Schrödinger, E. 1944. *What is Life?* Cambridge: Cambridge University Press.
- Schwartzman, D. 1999. *Life, Temperature and the Earth: The Self-organizing Biosphere*. New York: Columbia University Press.
- Sober, E., and D. S. Wilson. 1997. *Unto Others: The Evolution and Psychology of Unselfish Behaviour*. Cambridge: Harvard University Press.

- Staley, M. 2002. Darwinian selection leads to Gaia. *Journal of Theoretical Biology* 218 (1): 35–46.
- Stern, N. 2006. *The Economics of Climate Change: The Stern Review*. Cambridge: Cambridge University Press.
- Volk, T. 1998. *Gaia's Body—Toward a Physiology of the Earth*. New York: Copernicus.
- Volk, T. 2004. Gaia is life in a wasteworld of by-products. In S. H. Schneider, J. R. Miller, E. Crist, and P. J. Boston, eds., *Scientists Debate Gaia: The Next Century*. Cambridge: MIT Press, pp. 27–36.
- Watson, A. J. 2004. Gaia and observer self-selection. In S. H. Schneider, J. R. Miller, E. Crist, and P. J. Boston, eds., *Scientists Debate Gaia: The Next Century*. Cambridge: MIT Press, pp. 201–208.
- Watson, A. J., and J. E. Lovelock. 1983. Biological homeostasis of the global environment: The parable of Daisyworld. *Tellus* 35B: 284–89.
- Wilkinson, D. M. 1999. Is Gaia really conventional ecology? *Oikos* 84: 533–36.
- Williams, H. T. P., and T. M. Lenton. 2007a. The flask model: Emergence of nutrient-recycling microbial ecosystems and their disruption by environment-altering “rebel” organisms. *Oikos* 116 (7): 1087–1105.
- Williams, H. T. P., and T. M. Lenton. 2007b. Artificial selection of simulated microbial ecosystems. *Proceedings of the National Academy of Sciences* 104 (21): 8918–23.
- Williams, H. T. P., and T. M. Lenton. 2008. Environmental regulation in a network of simulated microbial ecosystems. *Proceedings of the National Academy of Sciences* 105: 10432–37.
- Wood, A. J., G. J. Ackland, J. Dyke, H. T. P. Williams, and T. M. Lenton. 2007. Daisyworld: A review. *Reviews of Geophysics* 46: RG1001.



-

-



F

