Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **533,** 2379–2390 (2024)
Advance Access publication 2024 August 20

https://doi.org/10.1093/mnras/stae2003

What does not kill *Gaia* makes her stronger: impacts of external perturbations on biosphere evolution

R. Arthur,[1] A. E. Nicholson[2]★ and N. J. Mayne[2]

[1]*Department of Computer Science, Faculty of Environment Science and Economy, University of Exeter, Exeter EX4 4QL, UK*
[2]*Department of Physics and Astronomy, Faculty of Environment Science and Economy, University of Exeter, Exeter EX4 4QL, UK*

Accepted 2024 August 15. Received 2024 August 14; in original form 2024 April 25

ABSTRACT

Life on Earth has experienced numerous upheavals over its approximately 4 billion year history. In previous work, we have discussed how interruptions to stability lead, on average, to increases in habitability over time, a tendency we called entropic *Gaia*. Here, we continue this exploration, working with the Tangled Nature Model of co-evolution, to understand how the evolutionary history of life is shaped by periods of acute environmental stress. We find that while these periods of stress pose a risk of complete extinction, they also create opportunities for evolutionary exploration which would otherwise be impossible, leading to more populous and stable states among the survivors than in alternative histories without a stress period. We also study how the duration, repetition and number of refugia into which life escapes during the perturbation affects the final outcome. The model results are discussed in relation to both Earth history and the search for alien life.

Key words: astrobiology – Earth – exoplanets – planets and satellites: detection.

1 INTRODUCTION

The long history of life on Earth is marked by a number of ‘revolutions’ (Lenton & Watson 2013), large changes in ecosystems and biogeochemical reaction networks. For example, the evolution of oxygenic photosynthesis fundamentally changed the surface chemistry of our planet, leading to the dramatic build up of oxygen in the atmosphere roughly 2.5 billion years ago, known as the great oxidation event. This build up of oxygen not only caused widespread poisoning to life that had not yet adapted to an oxygen-rich environment, but is also thought to have triggered a global glaciation event (Sahoo et al. 2012; Lenton & Watson 2013). Much later, the expansion of plants on land, significantly increased chemical weathering and therefore reduced atmospheric CO_2 (Lenton et al. 2012; Porada et al. 2016). This is hypothesized to have triggered a global glaciation and subsequent (Late Ordovician) mass extinction. Non-biotic perturbations such as changes in volcanism, asteroid impacts, plate tectonics, and climate change have also caused widespread extinctions and resulted in the emergence of new ecosystems (Bond & Grasby 2017). Even in cases where the ultimate cause is non-biotic, for example, asteroid impacts (Alvarez et al. 1980) or volcanism (Campbell et al. 1992), it is still often the case that life participates in positive feedback loops which worsen conditions and accelerate the extinction event (Bond & Grasby 2017; Dal Corso et al. 2022).

These periods of stress and their associated mass extinctions are thought to pose a problem for the body of work known as *Gaia* theory, which posits that life interacts with the non-living Earth so as to maintain and even improve conditions for life (Lovelock & Margulis 1974). Arguments against *Gaia* often contrast the stable biogeochemical cycles currently observed against runaway feedback enhanced by life before or during extinction events. There are also particular instances where some species (or group of species) acts to worsen conditions for itself (Kirchner 2002, 2003). Taken to the extreme, some have argued that life is in fact self-destructive (Ward 2009). The question of whether life is self-reinforcing or self-destructive clearly has profound implications not only for our understanding of Earth history but also on the search for life on other worlds. The answers will shape our expectations for the prevalence of life, especially complex life, in the universe. In particular, for exoplanets (planets outside the Solar system), a deeper understanding of the life–climate interaction and the impact of large-scale perturbations may well prove vital in selecting the few targets for which expensive and intensive follow-up biosignature observations are performed.

In previous work (Arthur & Nicholson 2017; Lenton et al. 2018; Arthur & Nicholson 2022, 2023a), we have taken the position that these revolutions are a feature, not a bug. Over the course of Earth history, there have been numerous cycles of extinction and recovery. We propose that these cycles should be considered as part of *Gaia*. In particular, they are the mechanism by which *Gaia* can make large jumps in complexity, where complexity is defined, broadly, to mean the number of processes and components in a connected system (Adami 2002; Gell-Mann 2002). We call this process sequential selection with memory and the mechanism of growth by extinction and recovery of the entropic ratchet (ER). This idea, which originated from studying a particular co-evolutionary model (Arthur & Nicholson 2022), is more general than that model. It also occurs in simple probability models (Arthur & Nicholson 2023a) and here we will argue that there is evidence for it in the record of life on Earth. The ER mechanism is implied by the following three conditions repeating over and over again:

★ E-mail: arwen.e.nicholson@gmail.com

(i) Gaian homeostasis can be destabilized by the evolution of new species or by exogenous events.

(ii) These events cause some, or all, of the ‘core’ species (that is, species which are key to the biogeochemical cycles that maintain the present conditions) of the global ecosystem to go extinct. New core species then arise, which create new niches and participate in biogeochemical cycles (Alroy 2008).

(iii) These extinctions are not total, for example, core species can survive but become too rare to fulfill their prior ecosystem function (Hull, Darroch & Erwin 2015). This means the post-event biodiversification, that fills the ecological space opened by the mass extinction, can build on evolutionary innovations of the preceding period, which tends to result a higher complexity and diversity of life.

Points 1 and 2 describe sequential selection (Lenton et al. 2018) – periodic resets of the global ecosystem. Adding 3 describes sequential selection with memory – periodic resets that retain remnants of previous global ecosystems and build on their evolutionary innovations. The repetition of 1, 2, and 3 is what we refer to as the ER, where the name is meant to evoke the idea that entropy, a combination of diversity and abundance (see Arthur & Nicholson 2022 for more precise definitions) changes in discrete jumps with each reset. Conditions 1 and 2 are likely uncontroversial. As well as distinguishing the ER mechanism from simpler ideas, point 3 is also why we claim this process is Gaian. Our *Gaia* is the entropic *Gaia* of Arthur & Nicholson (2017, 2022), referring to an evolutionary trend in the direction of greater biomass, species diversity and life-enhancing abiotic interactions and Gaian processes are those which enhance this trend.

After each extinction event, global biomass, and diversity could decrease or increase relative to the previous baseline, any particular event must be analysed and understood individually. However, as we have argued elsewhere (Arthur & Nicholson 2022, 2023a), there is a general tendency for systems with ‘memory’ to increase in complexity over time, a point often noted in complex systems theory (Anderson et al. 2004) and seen in other settings from probability models (Arthur & Nicholson 2023a) to glasses (Ritort 1995). This increasing complexity is often measured by the proxy of entropy (see e.g. Roach et al. 2017; Arthur & Nicholson 2022), hence the name ER. In a very abstract sense, as discussed in Arthur & Nicholson (2022), sucessive resets move the system to a ‘position’ in the space of ecosystems which has the potential to support more diversity and abundance. The system’s memory can be thought of as the current position on the ‘path’ through the space of possible ecosystems. More concretely, for life on Earth, this memory is the global biota, which is reduced but not eliminated by these resets. Evolutionary innovations, like photosynthesis or oxygenic metabolism, are not generated from scratch after every catastrophe. Our claim here is that repeated resets lead to a trend of increasing diversity and abundance together with a reduction in the rate of mass extinction, which makes the Earth (or any inhabited planet) ‘more Gaian’ over time.

In previous work, we have demonstrated this mechanism operating in a general model of co-evolution (Arthur & Nicholson 2017, 2022). In the rest of this section, we argue that this mechanism also has support from the literature on mass extinction and biodiversity change over geological time. For example, Benton (1995) and Newman & Sibani (1999) demonstrate a trend over the Phanerozoic (the last ~ 540 Mya) of increasing biodiversity, despite numerous mass extinctions. For particular mass extinction events, there is some evidence of increases in post-extinction complexity. Since complexity is associated with a number of ecological features: having multiple parts, diversity, interactions, emergence, non-linearity, and historicity (Elliott-Graves 2023), we assess its change through a range of proxies. For example, from the relative abundance distribution of marine fossils Wagner, Kosnik & Lidgard (2006) finds that complex ecosystems are more common during the Meso-Cenozoic (later Phanerozoic) than during the Paleozoic (earlier Phanerozoic) where the boundary is marked by Permian–Triassic mass extinction. Similar ideas have been discussed in the literature on mass extinction, notably the idea of Earth system succession – ‘the sequential change in global ecosystems that occurs as biogeochemical reservoirs and fluxes return to equilibrium after perturbation’ (Hull 2015). Our claim is that these extinction events, which greatly disrupt contemporary biota are, on average, positive for life in the long run, that is, over spans of time measured in 10s or 100s of millions of years which may incorporate a number of such events.

Since life has a profound effect on (bio)geochemical cycles, large-scale disruptions of global ecosystems can impact these cycles. The way in which life recovers after such events can be complex (Solé, Montoya & Erwin 2002; Solé et al. 2010) but these disruptions open the possibility for large changes in biogeochemical regulation that would not be possible without such events. To select some major examples

(i) The Great Oxidation Event caused tremendous damage to existing anaerobic species, which, for the most part, could not tolerate a high oxygen environment, and also precipitated a possible Snowball Earth period (Huronian glaciation). However, the evolution of Eukaryotic and multicellular life was enabled by the higher oxygen concentrations which allowed for aerobic respiration to become dominant resulting in far more available energy for life (Mills et al. 2022).

(ii) The Cryogenian/Snowball Earth period at the end of the Proterozoic, precipitated by the Neoproterozoic Oxygenation Event, would have certainly had a negative effect on extant life. Afterwards, we see recovery and (enabled by higher oxygen levels) the further complexification and diversification of life, culminating in the Ediacaran biota (Narbonne et al. 2012).

(iii) The causes of the end-Ediacaran extinction are uncertain, with explanations ranging from a standard mass extinction event to a more gradual biotic replacement (Laflamme et al. 2013; Darroch et al. 2018). Whatever the causes, the mass disappearance of the Ediacaran biota was immediately followed by the Cambrian explosion.

(iv) During the Phanerozoic there have been numerous mass extinction events, with scholarship mostly focusing on recovery over shorter time-scales. The end-Ordovician extinction is one example. Preceded by the ‘Great Ordovician Biodiversification Event’ (GOBE) and the colonization of land by plants (Lenton et al. 2012), the result was the increase of atmospheric oxygen and fire-mediated feedback to stabilize atmospheric oxygen concentrations at ~ 20 per cent (Lenton et al. 2016). This represents an increase in the complexity of the biogeochemical feedback network.

We do not claim that all mass extinction events in Earth history need have a positive impact on species diversity or abundance. For example, despite the findings of Wagner et al. (2006), diversity levels reached during the GOBE took tens of millions of years to recover after the largest mass extinction event of the Phanerozoic, the end-Permian (Raup & Sepkoski 1982; Rohde & Muller 2005). Other events are just ‘blips’ on *Gaian* time-scales. For example, the Cretaceous–Paleogene event, generally agreed to be caused by an asteroid impact (Alvarez et al. 1980), appears to have had little long-term impact on trends in biodiversity (Rohde & Muller 2005), extinction rate (Benton 1995) or global temperature (Scotese et al. 2021).

The complexity of the global ecosystem and life-environment feedback (i.e. *Gaia*) has increased over geological eons. While the consequences of complexity in ecology are debated (Landi et al. 2018), our view, supported by our models of co-evolving ecosystems (Arthur & Nicholson 2017, 2022, 2023a), is that an increase in complexity is associated with an increase in habitability and stability. In real ecologies complexity, measured via, for example, biodiversity, can enhance stability in a number of ways, from functional redundancy (Rosenfeld 2002) buffering against local extinctions to the stabilization of global biogeochemical cycles. An example of the latter is the evolution of land plants which increased the efficiency of silicate weathering. Their evolutionary adaptation to different climate, temperature, topography, etc. means that these factors exert less influence on the silicate weathering cycle (Lenton et al. 2016; Payne et al. 2020).

There is less work on how the total abundance of life has changed over time. Marine environments do show increased abundance over geological time (Bambach 1993; Martin 1996; Allmon & Martin 2014) and the evolution of plants resulted in enormous increases in the mass of Earth's biota (McMahon & Parnell 2018). Complex systems of recycling can also increase abundance, allowing limiting nutrients to 'go further' than would otherwise be possible, for example, phosphorus is the limiting nutrient in most aquatic ecosystems where recycling ratios of around 46:1 are quoted in Wilkinson (2023).

Diversity, abundance and stability should be part of any definition of 'planetary habitability'. According to our framework, increases in these features are ultimately caused by life interacting with life and the planet in a way that is conducive to life, even if disasterous in the short term, hence our identification of this mechanism as a Gaian process. *Gaia*, like life itself, should not be expected to have emerged *de novo*, fully formed and functional, nor to be eternally static. It is reasonable that *Gaia* can and should evolve.

Of more relevance to astrobiology than Earth history (where we know that life survived all previous extinction events) is the idea of selection by survival (SBS) (Bouchard 2014; Ford Doolittle 2014; Lenton et al. 2018; Arthur & Nicholson 2023a, b). This is the obvious fact that only those planets where life survives mass extinctions have life! This could mean that the surviving life on those planets has some special properties, or the events on the surviving planets were less severe. In an astrobiological context, SBS operates alongside the ER idea discussed above. Regular mass extinctions drive long term increases in habitability, but are also opportunities to lose the game entirely and end up with a planet devoid of life. For this reason, we study the interplay between SBS and the ER in our model ecosystems, and try to understand what this means for searches for extraterrestrial life.

The detection and subsequent characterization of exoplanets, planets orbiting stars other than the Sun, has provided a vast number of potential candidates for non-Earth biospheres. In fact, a major goal of the field of exoplanet research is to potentially detect a 'biosignature' (Catling et al. 2018) in the atmosphere of an exoplanet (e.g. the LIFE mission, Quanz et al. 2022). For this case, a biosignature must be detectable, requiring life to both have survived previous extinction events and established a large-scale interaction with the climate, i.e. have created an 'exo-*Gaia*' (Nicholson et al. 2018). Although the number of potential exoplanets is vast the resources required to perform detailed observations, modelling and analysis of a potential biosignature mean that targets will have to be carefully selected.

In this paper, we study the impacts of large exogenous perturbations in a model of a planetary ecosystem over geological time. Most of the great extinction events in Earth history are thought to have arisen from a combination of abiotic and biotic factors. In previous work (Arthur & Nicholson 2017, 2022), we only considered biotic effects, the disruption of a stable period by the evolution of new species. Here, we also introduce abiotic effects, captured as a disruption of a stable period by a sudden decrease in carrying capacity. In Section 2, we review the idea of refugia (a location supporting an isolated population during some period of environmental stress), and introduce our model in Section 3. We study, in detail, the effect of a single perturbation in Section 4. We extend this in Section 5 to look at the effect of the duration of the perturbation, the effect of repeated perturbations, and compare outcomes where there is one big refugium to a number of smaller ones. We conclude in Section 6 with a discussion of the implications of these results for Astrobiology.

2 REFUGIA

As discussed previously in this work and others (Arthur & Nicholson 2022) gradually increasing diversity is the key to increasing habitability and that this can be maintained by various means which we referred to as *Gaia's* memory. One form this memory takes is as *refugia*, areas of tolerable conditions amidst an uninhabitable environment. Refugia originally referred to the restricted ranges of various species during glacial maxima, particularly during the last ice age (Stewart & Lister 2001). This body of work studies and identifies macro-refugia (Ashcroft 2010), for example, lower latitudes that avoid glaciation and cryptic/micro-refugia (Stewart et al. 2010), for example, temperate areas within glacial zones. This is achieved through analysis of the pollen record (Bennett, Tzedakis & Willis 1991) or through genetic evidence (Cheddadi et al. 2006), since periods in a refugium usually correspond to a population bottleneck that leaves a distinct signature in the modern species' DNA.

The most extreme planetary glaciations, snowball Earths, are also posited to have had refugia. Although most research supports the existence of open ocean conditions at equatorial and lower latitudes (Hyde et al. 2000; Peltier et al. 2004; Song et al. 2023) which would act as refugia, some models (Braun et al. 2022) predict a hard snowball. In such cases, microrefugia are still a possibility. Notable examples of such microrefugia are suggested by Campbell, Waddington & Warren (2011, 2014), who claim that narrow seas (like the modern Red sea) could provide refugia for photosynthetic eukaryotic algae. Windblown dust can lower the albedo of glaciers (Abbot & Pierrehumbert 2010) which can lead to pockets of liquid water called cryoconite holes. The same process acting on a larger scale on snowball earth could have created refugia for eukaryotes tolerant of cold water, low salinity, and strong radiation (Hoffman 2016). Similarly, areas of 'dirty ice' (Hawes et al. 2018) could have provided stable and nutrient-rich micro-refugia, as they do presently in Antarctica. Lechte et al. (2019) propose that the mixing zones of oxygen-rich glacial meltwater with iron-rich seawater could have provided sufficient energy for chemosynthesis and therefore represent another type of refugium. Hot springs have also been proposed as providing 'Noah's Arks' for photosynthetic life during hard snowball Earth events (Schrag & Hoffman 2001; Costas, Flores-Moya & López-Rodas 2008).

A number of different types of refugia are described by Bennett & Provan (2008), demonstrating a diversity of ways that species can survive periods of climatic stress by altering their abundance and distribution. Of most relevance for us are the classical and tropical refugia, when species restrict their range to one (classical) or many (tropical) small areas in an otherwise inhospitable environment. We also note that a refugium is usually species specific (Stewart et al. 2010) and so a refugium for one species may be inhospitable for another. This work does not seek to model any particular glaciation or other event from Earth history. Rather we aim to study how, in general, retreat to refugia affects the long-term habitability and hence chance of life detection on any planet. Therefore, we adopt a somewhat broader concept of refugia (Keppel et al. 2012) applying the concept to whole ecosystems, having in mind something like the open equatorial ocean or habitable narrow seas during snowball Earth events.

We also note that while there has been much work on the recovery of diversity after mass extinctions (e.g. Condamine, Rolland & Morlon 2013), there has been less emphasis on the recovery of biomass or productivity, reflecting a general emphasis in ecology which tends to be more interested in diversity than abundance (Wilkinson 2023). However, the role of abundance is crucial for understanding ecosystem function (Spaak et al. 2017). Indeed, Hull et al. (2015) make the interesting point that it is not strictly necessary for a species to go extinct to trigger a mass extinction, rather core species, groups of organisms which participate in key biogeochemical cycles, can fall below the abundance threshold required for them to effectively perform their roles in the cycle, see also Avolio et al. (2019). Biosphere abundance is also particularly important when searching for signs of life on distant planets; in order for life to be remotely detectable it must exist in sufficient quantifies to influence its planet in a significant way (Seager, Bains & Hu 2013).

In summary, refugia represent a way for life to survive during inhospitable conditions. At least some refugia are necessary so that life as a whole does not die out. These refugia act as a memory and storehouse of genetic diversity for *Gaia*. We will therefore be interested in understanding how the number and type of refugia interact with the ER and SBS effects and what this means for the probability of complex life on a planet.

3 THE TANGLED NATURE MODEL

The Tangled Nature Model (TNM) (Christensen et al. 2002; Jensen 2018) is a framework for understanding co-evolving species. The TNM is characterized by periods of stability where groups of species persist for a significant time, often called a *quasi-Evolutionary Stable Strategy* or qESS. These qESSs are interrupted by ‘quakes’, where the ‘core’ of the species network is disturbed by a newly evolved species and is rearranged or collapses completely. After a quake, the system finds a new qESS and the total population and composition of the biosphere drastically changes. These quakes are an inherent feature of the TNM and require no external perturbation.

One of the key characteristics of the TNM is the tendency for the biosphere to increase in total population, diversity, and stability over time. At later times, the TNM biosphere is more robust and less prone to quakes and thus the periods of stability get longer as time goes on. The quakes in the model are not a deterrent to this increasing stability but rather the mechanism by which this is achieved. Previous work (Piovani, Grujić & Jensen 2016; Arthur & Nicholson 2017) has demonstrated how the TNM model is closely related to the logistic model of population dynamics. Thus, because the TNM arises from consideration of very general principles it is reasonable to posit that the model results are also of wide applicability.

In the original formulation of the TNM, a parameter μ represents the ‘abiotic’ carrying capacity and remains constant throughout experiments, while the growth rate of species within the biosphere depends on μ as well as the other species extant at that time. Later work (Arthur et al. 2017) allowed species to directly impact the carrying capacity of the system and demonstrated that species-environment co-evolution leads to TNM biospheres tending to increase the abiotic carrying capacity over time.

Table 1. TNM parameter values.

C	100
θ_J	0.25
σ	0.01
θ_K	0
μ_0	0.1
ν	5×10^{-6}
p_k	0.2
p_{mut}	0.01
L	20

The TNM has been described numerous times in great detail (Christensen et al. 2002; Arthur et al. 2017; Arthur & Nicholson 2022) and the reader is referred to these works for a more detailed specification of the model. Briefly, species i are labelled by a length L binary genome. The population of species i is N_i and the total population of all species is $N = \sum_i N_i$. Each species has a fitness, f_i, which depends on the other extant species in the model given by

$$f_i = C \sum_j J_{ij} \frac{N_j}{N} + \sigma \sum_j K_{ij} N_j - \mu_0 N - \nu N^2, \tag{1}$$

where C, σ, μ_0, ν are constants and the sums are over all extant species. J_{ij} is a matrix of direct interspecies interactions, K_{ij} is a matrix of species environment interactions (the effect of j on the environment of i). The values of J and K are chosen at random from a standard normal product distribution used for reasons of computational convenience (Arthur et al. 2017) where a fraction of the entries, $1 - \theta_J$ and $1 - \theta_K$, respectively, are set to 0. μ_0 is the reciprocal of the total carrying capacity, while ν is a very small damping factor which is irrelevant except in very rare cases of extremely high populations. Setting $\sigma = 0$, $\nu = 0$ corresponds to the original TNM of Christensen et al. (2002), non-zero values give the version where species affect the environment proposed in Arthur & Nicholson (2022).

The model consists of repeating the following steps

(i) select an individual at random and kill it with probability p_k; and

(ii) select an individual and reproduce it with probability $p(f_i) = \frac{1}{1+\mathrm{e}^{-f_i}}$

Each reproduction copies the binary genome of the individual with probability p_{mut} to flip one of the digits, potentially creating a new species. The time-scale for the model is measured in ‘generations’ which consist of N/p_k repetitions of the two steps above. We use standard values for the various parameters, listed in Table 1.

In summary, the TNM, as employed in this work, tracks populations of species which interact with each other and their environment with each timestep providing random death, reproduction and mutation of individuals. In general, TNM systems evolve to higher complexity (here indicated by the increase in the population) and higher stability, moving through increasing stable periods (qESS) disrupted when the core of the life network is disrupted through the natural evolution of destabilizing species (quakes). Concepts such as habitability, stability, and complexity can be made very concrete and measurable in this framework, see Arthur & Nicholson (2017, 2022, 2023b). We refer interested readers to those works for details of how the ER mechanism in the TNM interacts with these concepts. We

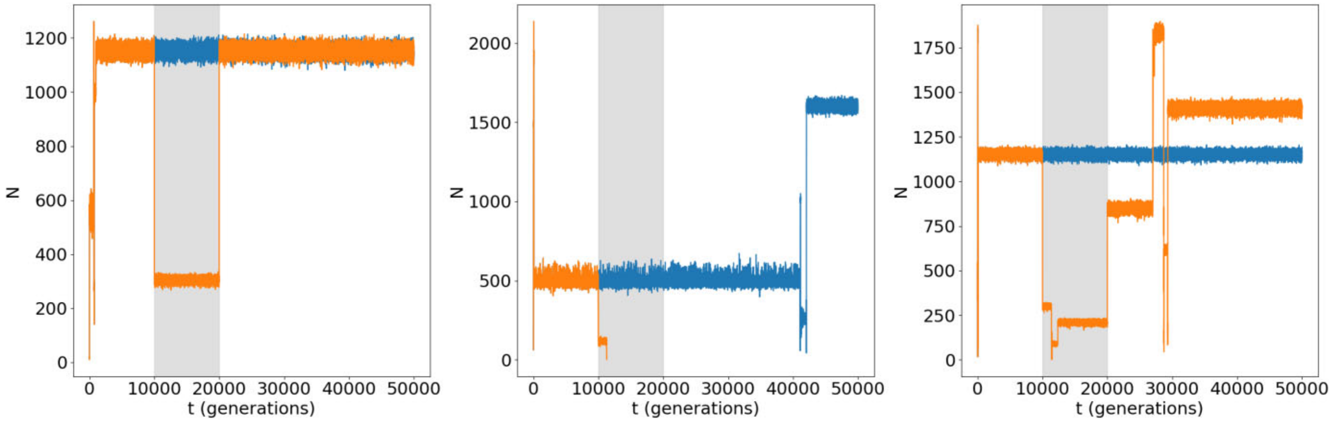

Figure 1. Some examples of TNM runs where we compare cases without perturbation (blue) to cases with perturbation $\mu_0 \rightarrow \mu = 0.4$ (orange). The shaded grey area is where the orange runs experience the perturbation. From the left to right panels, we see minimal effect, perturbation induced extinction, and perturbation induced divergence.

will mostly interrogate the model by looking at population, which is closely related to number of core species, total diversity, number of quakes and entropy.

The purpose of this work is to understand the effect of an external perturbation which suddenly reduces the carrying capacity, representing some change in the environmental context, where life is forced to survive in one or multiple refugia. We will study the effects of these perturbations by looking at changes in total abundance, in individual runs or in averages over many runs. Interested readers are referred to our previous work to see how other variables change on long time-scales. In the experiments, we perform in this work, the other macroscopic TNM variables have similar behaviour to the total population. Perturbations are implemented by abruptly changing the value of $\mu_0 \rightarrow \mu$ for a set number of generations and then reverting to μ_0 when the stress period is over. All other parameters remain fixed during the perturbation. Higher values of μ correspond to worse conditions for life and vice versa.

4 ONE REFUGIUM

Our first experiment is similar to one preformed, in a very different context, by Arthur et al. (2017). We allow the model to run as normal for 10^4 generations then abruptly increase the value of μ, run the model at that value for 10^4 generations, then reset it to the original value and continue for another 3×10^4 generations. Some illustrative examples of different model runs are shown in Fig. 1. We can conceptualize this period of lower μ as something like Snowball Earth, where the planet can support a much lower abundance of life which persists in a refugium. The examples in Fig. 1 have been chosen deliberately to illustrate the most important possible results of a perturbation: no long-term effect, perturbation-induced extinction, and perturbation-induced divergence.

Fig. 2 summarizes 1000 realizations of the model, for one particular value of the perturbation $\mu = 0.4$. The number of runs experiencing complete extinction of all individuals increases due to the perturbation. Fig. 2 also shows the average population of the subset of runs which survive to the end of the perturbed and unperturbed experiments. The average population increases logarithmically with time. Notably, after the perturbation, there is a jump in the average population above the unperturbed baseline which persists until the end of the run. The fact that bad conditions make total extinction more likely is quite intuitive. What is less intuitive is that runs which survive are 'better off' when they have experienced a perturbation than they otherwise would have been.

Fig. 3 shows the number of surviving runs as well as the difference in population between perturbed and unperturbed models at $t = 5 \times 10^4$ as a function of μ, the perturbation size. Expectedly, as the perturbation gets stronger, more runs experience total extinction. Less expectedly, the surviving runs where there has been a perturbation have higher final populations than ones which do not. There also seems to be a peak in the response to perturbation, with a maximum around $\mu = 0.4$. Note that since the average population increases approximately logarithmically (Becker & Sibani 2014), i.e. very slowly, a 5 per cent increase in population is quite significant and represents a leap forward by many thousands of generations.

To understand what is happening first, we note (Becker & Sibani 2014; Arthur & Nicholson 2022) that for a mutant species, a, to disrupt a qESS requires it to have high enough fitness to have significant reproduction probability, i.e. the species fitness should be above a minimum value ($f_{\min}$) set by

$$f_a > \log\left(\frac{p_k}{1-p_k}\right) = f_{\min}.$$

Using equation (1) this means

$$C\sum_j J_{aj}\frac{N_j}{N} + \sigma\sum_j K_{aj}N_j > f_{\min} + \mu N + \nu N^2. \tag{2}$$

We have set σ and ν to be quite small and the main requirement is that the new species growth rate, $r_a = C\sum_j J_{aj}\frac{N_j}{N}$ is large enough to overcome the 'barrier', on the right-hand side of equation (2), which is primarily set by the value of μN. In a qESS

$$\frac{\mathrm{d}N}{\mathrm{d}t} \simeq 0.$$

Using the mean-field approximation from Arthur & Nicholson (2022), and neglecting σ and ν, gives

$$N \simeq \frac{r}{\mu}, \tag{3}$$

for the population in equilibrium where

$$r = C\sum_{ij}\frac{N_i}{N}J_{ij}\frac{N_j}{N}.$$

A sudden increase in μ will not directly affect the species composition of the TNM system, so the value of r will be roughly the same

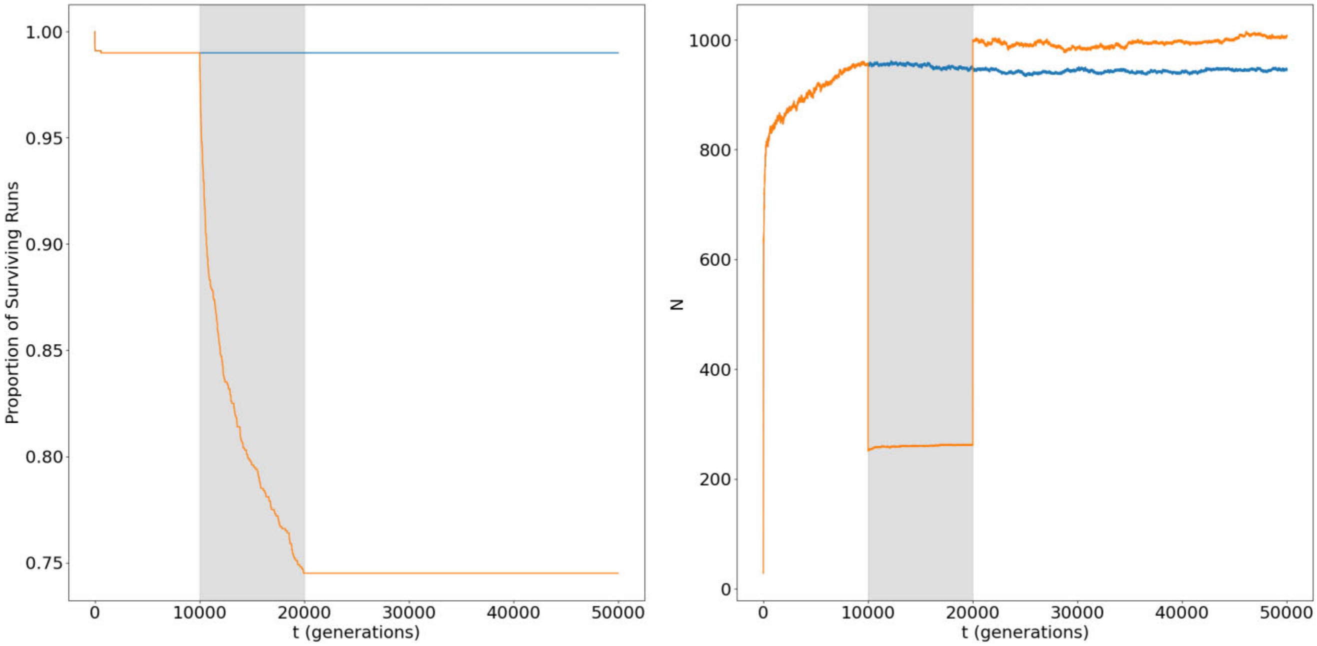

Figure 2. Averages without perturbation (blue) and with perturbation $\mu_0 \rightarrow \mu = 0.4$ (orange). The shaded grey area is where the orange runs experience the perturbation. Left panel shows the proportion of the $N = 1000$ runs which have survived up to that point. Right panel shows the average population *of the runs which survive the whole experiment*.

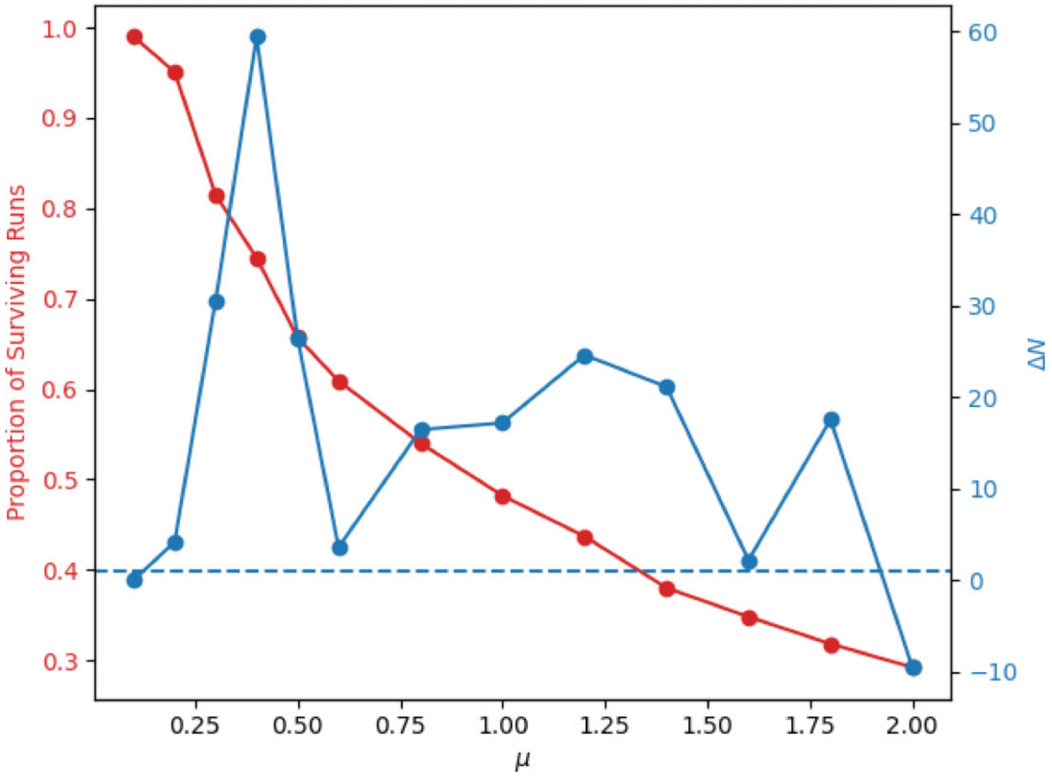

Figure 3. Proportion of runs which survive the whole experiment and ratio of population at the end of the surviving runs, with and without perturbation. Population of each run is measured by averaging the last 500 generations. Left (red) axis is the proportion of surviving runs, right (blue) axis is the population difference.

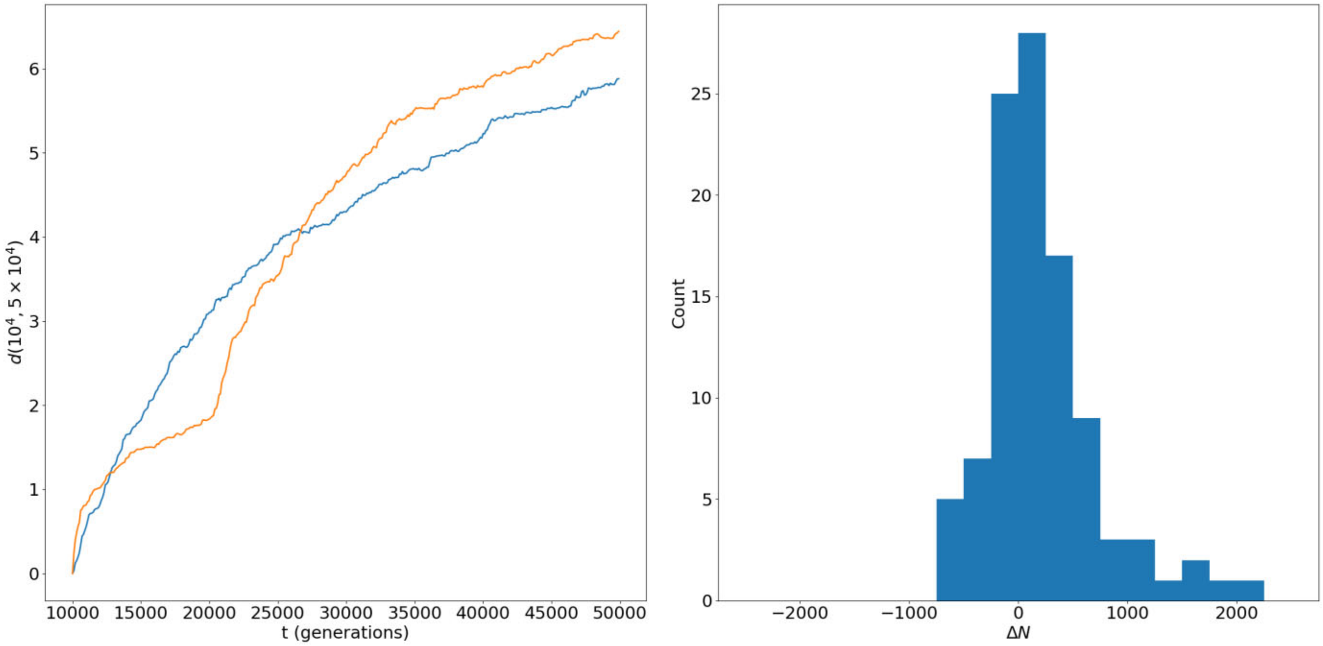

Figure 4. Showing the $\mu = 0.4$ perturbation. Left panel shows the 'genetic distance' between the core just before the perturbation and at the end of the run with (orange) and without (blue) a perturbation. Right panel shows the population excess among survivors of the perturbed case *when there has been a quake during the perturbation*.

immediately after the perturbation. Simply, the increase in μ will be compensated for by a decrease in N, while keeping the relative proportions of each species almost unchanged and thus leaving the barrier height r unchanged. Given the reduced rate of reproduction, one might therefore expect *fewer* quakes and since, as argued above, quakes are what drives the TNM to better (higher N) equilibria, the results of Figs 2 and 3 are at first glance puzzling.

Fig. 4 shows the genetic 'distance' between cores at $t_1 = 10^4$ (just before the perturbation) and $t_f = 5\times10^4$, measured by

$$d(t_1, t_f) = \sum_i \min_j H(\mathcal{S}(t_1)_i, \mathcal{S}(t_f)_j), \tag{4}$$

where $\mathcal{S}(t)$ refers to the set of core species' genomes at generation t and H is Hamming distance. d measures the smallest number of mutations required to get from the core at t_1 to the core at t_f. The perturbed curve (orange) shows some interesting features. Just after the perturbation, the distance increases – corresponding to core re-arrangement. After this initial jump, the rate of change decreases for the rest of the perturbation. Then for the $\sim 10^4$ generations after the perturbation, the rate of divergence rapidly increases. This rate of increase is enough to catch up and overtake the unperturbed systems so that, by the end of the experiment the perturbed systems manage to explore more of the 'landscape' (Arthur & Sibani 2017) and thus reach better final states.

Fig. 4, also shows a histogram of the population difference $\Delta N = N(5\times10^4) - N(10^4)$ for the surviving runs of the $\mu = 0.4$ perturbation *when there is a quake during the perturbation period in the perturbed run and not otherwise*. Operationally, a quake is defined as any change in the core composition together with at least a 5 per cent change in the total population N. The key point about this plot is that is is skewed right, towards higher population excess. This means, when there is a quake or core rearrangement during the perturbation the final population ends up higher than if there was no quake. The other possibilities: quake during the unperturbed and not the perturbed run, both quake and no quake, give symmetric distributions. This means it is runs which have perturbation induced quakes which are responsible for the increase in average population.

Close inspection of the runs which quake during the perturbation, and are responsible for the positive value of ΔN, indicate that the main cause of these is the variance in core species populations. While the population ratios N_i/N are fixed, the absolute populations are much smaller in the perturbation. In equilibrium we have $p(f_i) \simeq p_k$ and the expected number of reproductions of species i in one generation (N/p_k trials) is just the binomial expectation N_i. The binomial variance is $N_i(1 - p_k)$, and the square root of this measures the average fluctuation size. The signal-to-noise ratio is then $\sim \sqrt{N_i}$ i.e. when the population is low the variance around the mean value is, relatively, much higher.

This higher variance in N_i can cause spontaneous core-collapse when a fluctuation takes one of the core species to $N_i = 0$. However, the most significant effect observed is the translation of fluctuations in N_i to fluctuations in r which can, transiently but significantly, reduce the quake barrier. This makes a quake much more likely to happen. During a quake, the populations can get very low for a brief period (see e.g. Fig. 1), making the runs much more likely to go extinct. If a run survives the perturbation but at a worse (lower N) qESS it is also much more likely to go extinct. Quakes which cause increases in N are much more likely to survive, and therefore have higher r and so are more stable – this explains the initial jump and then plateau in d in Fig. 4. Runs which do not quake during the perturbation 'catch up' with the unperturbed runs after the perturbation goes away and the rate of reproduction increases, which explains the increase in d after the perturbation ends.

In summary: hostile conditions increase the importance of population fluctuations. This enables more and different quakes – allowing more exploration and enhancing the ER mechanism. However, quakes are risky, and only those which have positive outcomes survive to be counted – this is SBS. This could well have important

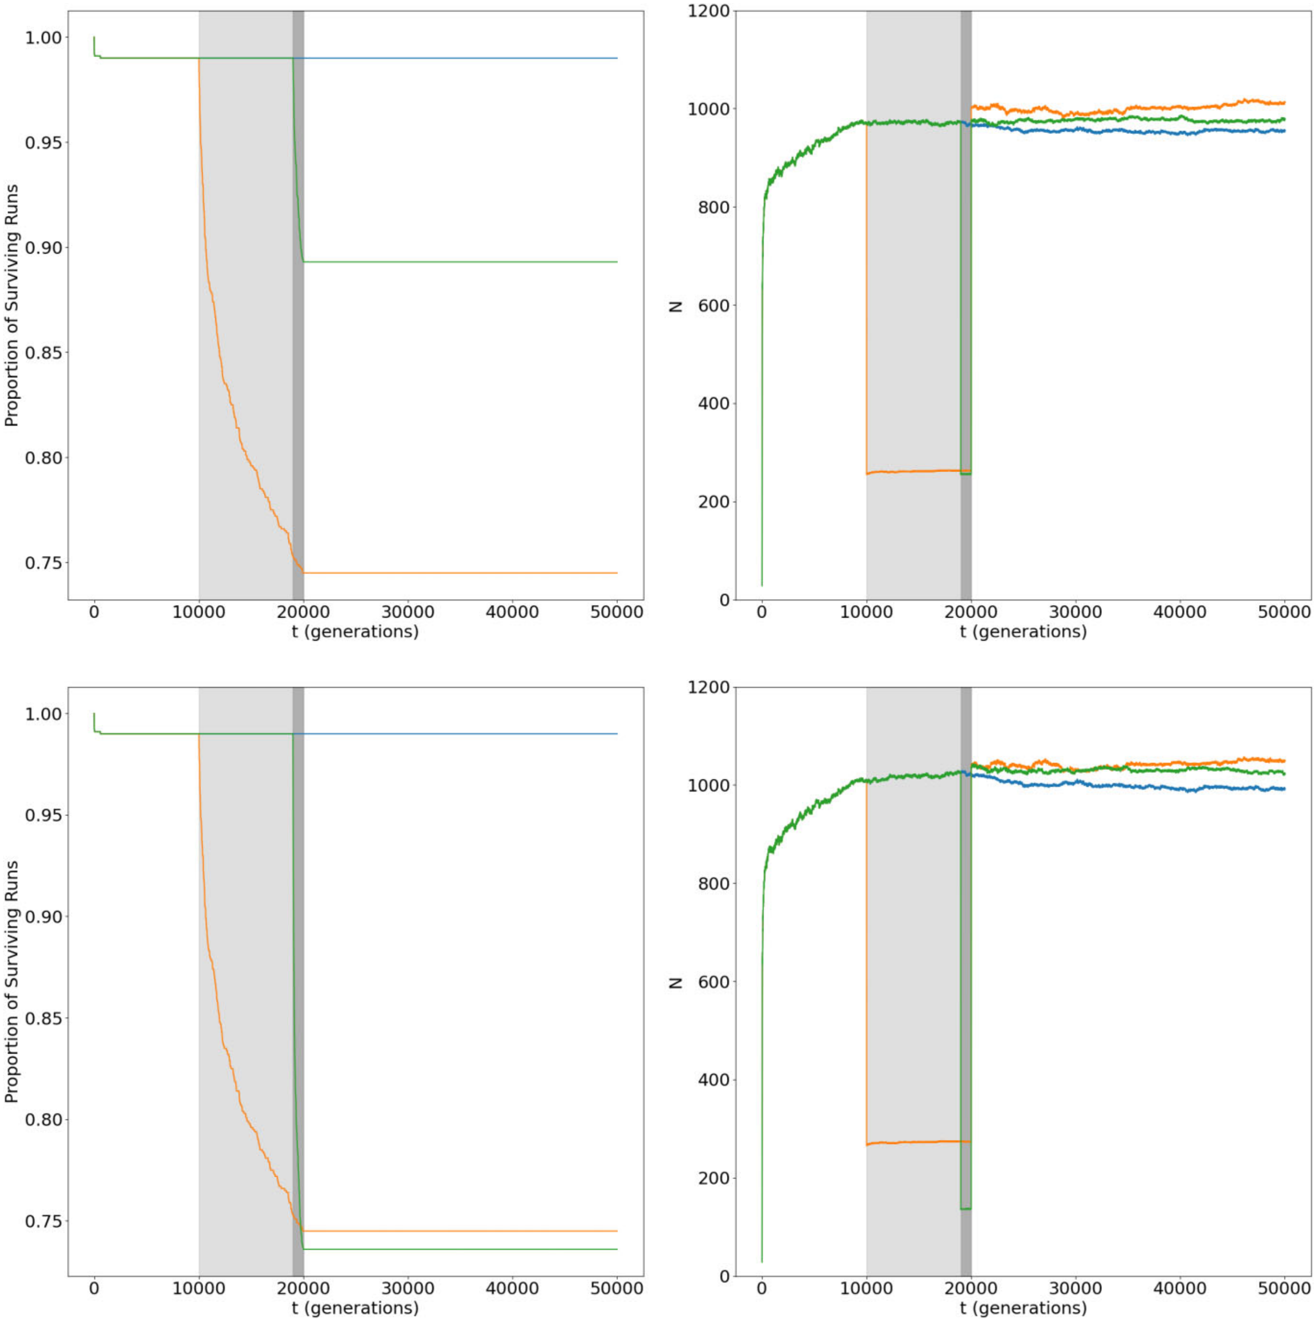

Figure 5. Showing long and short-duration perturbations. Only runs which survive all three conditions are included in the average. Blue: no perturbation, Orange: long perturbation, and Green: short perturbation. Left panel: proportion of survivors. Right panel: average population. Top panel: $\mu_{\text{long}} \rightarrow 0.4$ and $\mu_{\text{short}} \rightarrow 0.4$, Bottom panel: $\mu_{\text{long}} \rightarrow 0.4$ and $\mu_{\text{short}} \rightarrow 0.8$.

implications for our search for an inhabited exoplanet and merits further study.

5 OTHER PERTURBATION EXPERIMENTS

As discussed in the Introduction, ‘perturbations’ of various sizes, and severity have affected life on Earth. We do not seek to model any exactly, our model is too conceptual to simulate Earth history, however, we can explore some illustrative examples to give us an idea of how the above results are affected by the characteristics of the perturbation and life’s response to it.

5.1 Short perturbations

Fig. 5 shows the average of 1000 simulations performed as in the previous section along with some new simulations where the perturbations begin later, at $t = 19\,000$, and only last for 1000 generations. When the short perturbation is of the same severity, μ, as the long one, the effect on survival and on ΔN is less. However, when doubling μ so that the rate of extinction is similar, the effect on ΔN increases and the two cases are roughly equal. One could try to study the exact dependence of ΔN on μ; however, this depends on μ_0 and a number of other model parameters. The results above are

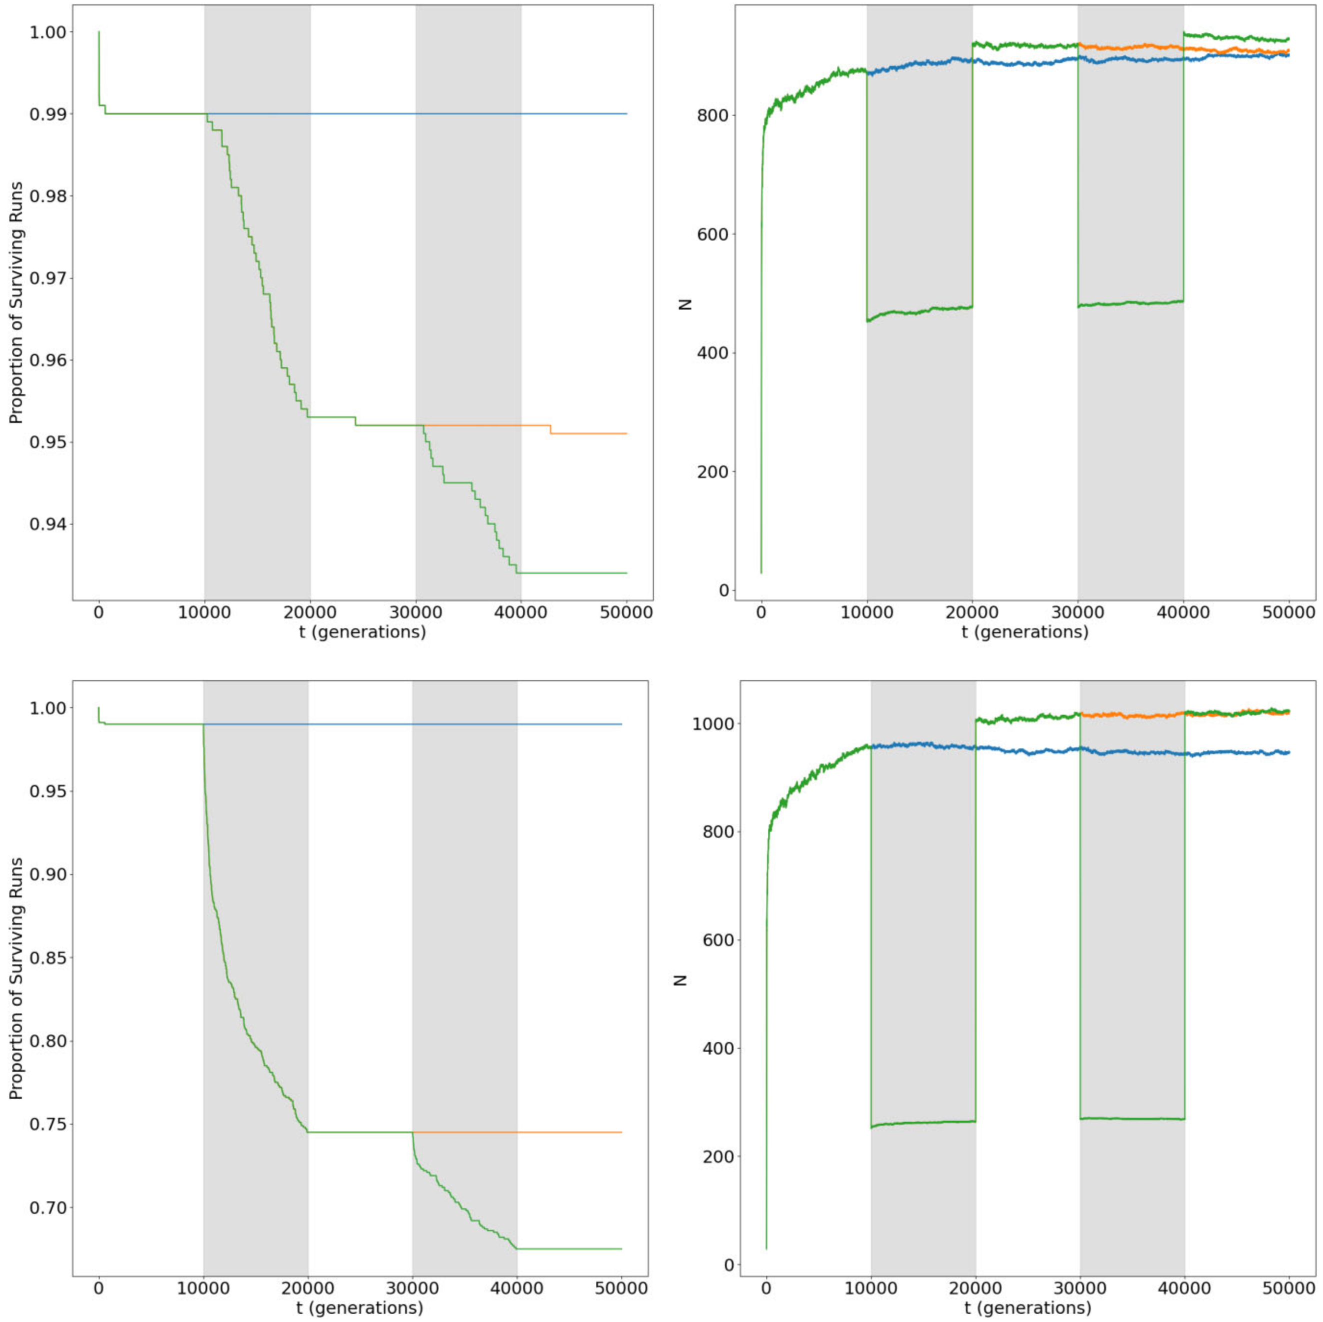

Figure 6. Showing 0, 1, and 2 perturbations. Only runs which survive all three conditions are included in the average. Blue: no perturbation, Orange: 1 perturbation, Green: 2 perturbations. Left panel: proportion of survivors. Right panel: average population. Top panel: $\mu \rightarrow 0.2$. Bottom panel: $\mu \rightarrow 0.4$.

illustrative of the fact that it is the total 'intensity' of the perturbation, duration × severity, that is key.

5.2 Multiple perturbations

Fig. 6 shows the effect of a second period of perturbation after the first. We see that for weak perturbations we can get a compounding effect on ΔN. However, for stronger perturbations the second period does not have any effect on ΔN. In the former case, the second period of perturbation seems to simply increase in perturbation intensity, by effectively increasing the duration. In the latter case, the initial perturbation seems to be sufficiently intense (that is severe enough and long enough) to have caused either a jump to a better and more stable qESS or a total extinction. The second perturbation is then acting on systems which have already been selected at this level and so has little effect, beyond some additional SBS ending more of the runs.

5.3 Multiple refugia

As discussed in Section 2, there are numerous ways in which species attempt to survive periods of stress. We have been studying the case where there is only a single refugia which is hospitable to life during the perturbation. It is also possible that multiple refugia exist and it

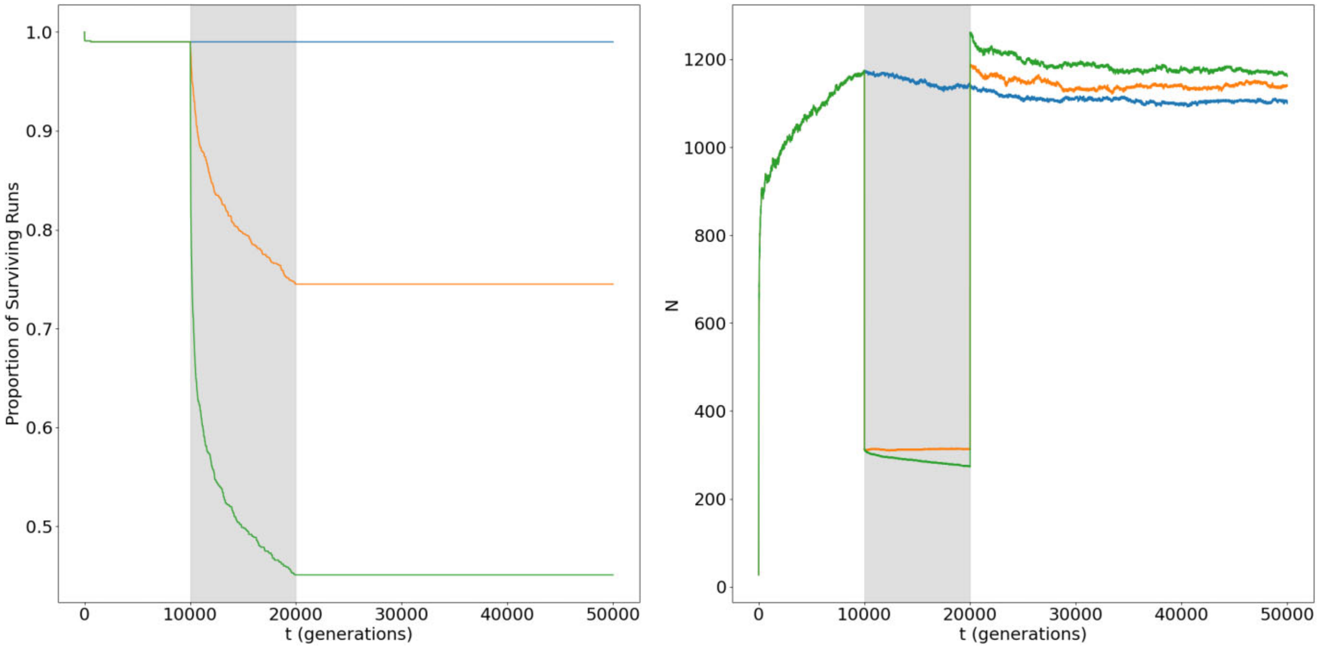

Figure 7. $K = 1$ versus $K = 4$ refugia. Blue: No perturbation, Orange: 1 refugia $\mu = 0.4$, and Green: 4 refugia $\mu = 1.6$. Left panel: proportion of survivors. Right panel: average population.

is interesting to ask what the difference would be between say, an open equatorial ocean (single refugia) or a number of open seas or habitable 'patches' of ice (multiple refugia) during a snowball Earth.

To simulate this, we run the TNM for 10^4 generations, then randomly distribute the individual agents in the system into K refugia. We then run these K systems independently for 10^4 generations with the perturbation μ. After this, we merge all surviving individuals from each of the K refugia back into a single system and run for another 3×10^4 generations at μ_0. Higher K at the same value of μ gives more chance for runs to survive, quake and so on. Using equation (3), $N \simeq \frac{r}{\mu}$, we can ensure the single refugia case has the same total carrying capacity as the K refugia case by using the value $K\mu$ during the perturbation for K refugia when we use the value μ for one refugia. We are asking if it is better (in terms of survival probability and ΔN) for an ecosystem to have all of its individuals contained within one big refugia or divided into many small ones, during external perturbations.

Fig. 7 shows one refugia at $\mu = 0.4$ compared to four at $\mu = 1.6$. We see that both the rate of extinction and the increase in ΔN are increased in the $K = 4$ case. $\mu = 1.6$ is quite an extreme perturbation and so many of the runs go extinct, even with $K = 4$ chances to survive the perturbation (i.e. four separate refugia). This extreme value of μ results in lower populations in each of the refugia, and thereby even greater chances of a quake occurring in any of them. Therefore, for runs where at least one of the refugia is inhabited at the end of the perturbation, the increased rate of quakes yield a higher population in the subsequent qESS state.

6 DISCUSSION

In this paper, we claim that the ER mechanism is a subtle but important influence on Life's trajectory over long time-scales. We have argued in the introduction that there is evidence from Earth history of this mechanism operating over geological time. In previous work, we have shown in great detail how the ER arises in the TNM (Arthur & Nicholson 2022), as well as in simpler models (Arthur & Nicholson 2023a). We call this a Gaian mechanism, in the specific sense defined in the introduction, a trend over long time-scales for Life to evolve in a 'direction' that ultimately supports more life. This mechanism depends on repeated cycles of crisis and recovery. Previous work on the TNM has shown precisely how the model spontaneously generates quakes and how this process leads to greater abundance, diversity and stability (Christensen et al. 2002; Becker & Sibani 2014). The similar cycles of crisis and recovery in Earth history discussed in the introduction suggest that it is plausible that this mechanism is operating on Earth over geological time-scales

This paper builds on previous work and aims to understand the effect of external perturbations on Life history, in particular, how they affect and interact with the ER mechanism and the spontaneous quakes that are characteristic of the TNM. We find that the reduced populations during the perturbation period allow for quakes to occur which would otherwise have been impossible. These either wipe out all life, or, allow 'long jumps' (Kauffman & Levin 1987) across the ecological landscape which tend to result in greater species abundance, and all of the consequences that abundance implies in the TNM, including higher species diversity and qESS stability, see Arthur & Nicholson (2022). These results should be interpreted carefully. It is true that *surviving* runs which endured the perturbation tend to have higher populations. From the perspective of the surviving population, the perturbations are ultimately helpful. However, many result in complete extinction. The idea that large events which are detrimental to carrying capacity might be harmful to life is expected, but the idea that they could be beneficial is not. Large perturbations present both an opportunity and a risk. By weakening the core, new possibilities are opened, at the cost of a significant risk of total extinction.

In terms of the selection mechanisms discussed in the introduction – both SBS and ER mechanisms are enhanced. SBS simply says any runs which survived the perturbation had to have properties which enabled their survival. In this case, they are runs which have higher than average abundance. More subtle, but potentially more interesting, is the fact that runs which survive tend to be better (having higher populations and the features associated with this in the TNM: higher diversity, core size, and reduced likelihood to quake), *because of* the perturbation enhancement to the ER.

These effects are really apparent when averaging over many possible life histories. Earth history is only a single time-series. There are suggestions, as noted in the Introduction, that large changes in the Earth System are often observed soon after a large perturbation. There is also the opposite observation, the so-called boring billion during the Proterozoic (Lenton & Watson 2013) was a period of relative stability and slow evolutionary innovation. Thus, these ideas have some support in Earth history, and reproducing these mechanisms in such a simplified framework allows us to understand the potential behaviour of the Earth over its deep past and future.

However, where this framework could prove vital is in the application for our search for habitable or indeed inhabited worlds beyond the Solar system. Many studies of exoplanets have been performed focused on identifying potentially 'habitable' planets through the application of 'abiotic' climate models, i.e. neglecting any potential life's impact on the climate. These have also, largely, focused on the modern Earth system (e.g. Fauchez et al. 2022) although work has begun recognising that perhaps a greater probable state would be that of the Archean Earth with its more simple biosphere (Arney et al. 2016). However, as discussed in this work, Life on Earth has had a huge impact on the climate (e.g. Lenton & Watson 2013) and it might be possible that habitable conditions can only persist for long time-scales on inhabited planets (Goldblatt 2016). Of course, modelling the complex interactions of a distant planetary climate system, including biogeochemical feedback from potential life forms is a significant challenge. However, as we detect more and more planets which are designated as potentially habitable we must begin to confront this problem and guide what will be resource-intensive follow-up observations to regions of exoplanetary parameter space that we deem most likely to host life. In this regard, simple model frameworks, as independent as possible of the nature of the system itself, are a powerful tool in beginning to map out this likelihood space. With the many thousands of potentially habitable exoplanets likely to exist in our local region of the galaxy alone, it is vital that we attempt to develop a statistical understanding of where we are *most likely* to find life. Our study suggests that although the probability of extinctions is enhanced by perturbations, for those systems that survive the diversity, complexity, and abundance of life are all increased over those which do not experience a perturbation. However, to extend this work to more specific predictions of where life might be most probable requires extensive follow-up and moving beyond conceptual models like the TNM. Previous work has explored interaction between life and the planetary climate to understand what determines the 'strength' of potential biosignatures (Nicholson et al. 2022) and how the biota influences its planet's climate (Nicholson & Mayne 2023). Taken together, these previous works and this study are building towards developing a probabilistic Gaian habitable zone where precious observational resources can be guided to where most fruitfully to undertake characterization follow-up campaigns designed to detect potential biosignatures.

ACKNOWLEDGEMENTS

This work was partly funded by the Leverhulme Trust through a research project grant [RPG-2020-82] and a UKRI Future Leaders Fellowship [MR/T040866/1].

DATA AVAILABILITY

Code used to generate data is available upon reasonable request from the authors

REFERENCES

Abbot D. S., Pierrehumbert R. T., 2010, J. Geophys. Res. Atmos., 115
Adami C., 2002, BioEssays, 24, 1085
Allmon W. D., Martin R. E., 2014, Paleobiology, 40, 256
Alroy J., 2008, Proc. Natl. Acad. Sci, 105, 11536
Alvarez L. W., Alvarez W., Asaro F., Michel H. V., 1980, Science, 208, 1095
Anderson P. E., Jensen H. J., Oliveira L., Sibani P., 2004, Complexity, 10, 49
Arney G. et al., 2016, Astrobiology, 16, 873
Arthur R., Nicholson A., 2017, J. Theor. Biol., 430, 177
Arthur R., Nicholson A., 2022, J. Theor. Biol., 533, 110940
Arthur R., Nicholson A., 2023a, AstroBiology, 23, 1238
Arthur R., Nicholson A., 2023b, MNRAS, 521, 690
Arthur R., Sibani P., 2017, Physica A: Statistical Mech. Appl., 471, 696
Arthur R., Nicholson A., Sibani P., Christensen M., 2017, Comput. Math. Organ. Theory, 23, 1
Ashcroft M. B., 2010, Journal of Biogeography, 37, 1407
Avolio M. L., Forrestel E. J., Chang C. C., La Pierre K. J., Burghardt K. T., Smith M. D., 2019, New Phytol., 223, 1106
Bambach R. K., 1993, Paleobiology, 19, 372
Becker N., Sibani P., 2014, Europhys. Lett., 105, 18005
Bennett K. D., Provan J., 2008, Quat. Sci. Rev., 27, 2449
Bennett K., Tzedakis P., Willis K., 1991, J. Biogeogr., 18, 103
Benton M. J., 1995, Science, 268, 52
Bond D. P., Grasby S. E., 2017, Palaeogeogr. Palaeoclimatol. Palaeoecol., 478, 3
Bouchard F., 2014, Biol. Theory, 9, 382
Braun C., Hörner J., Voigt A., Pinto J. G., 2022, Nat. Geoscience, 15, 489
Campbell I., Czamanske G., Fedorenko V., Hill R., Stepanov V., 1992, Science, 258, 1760
Campbell A. J., Waddington E. D., Warren S. G., 2011, Geophys. Res. Lett., 38
Campbell A. J., Waddington E. D., Warren S. G., 2014, J. Geophys. Res. Oceans, 119, 2679
Catling D. C. et al., 2018, Astrobiology, 18, 709
Cheddadi R. et al., 2006, Glob. Ecol. Biogeography, 15, 271
Christensen K., Di Collobiano S. A., Hall M., Jensen H. J., 2002, J. Theor. Biol., 216, 73
Condamine F. L., Rolland J., Morlon H., 2013, Ecol. Lett., 16, 72
Costas E., Flores-Moya A., López-Rodas V., 2008, New Phytologist, 180, 922
Dal Corso J. et al., 2022, Nat. Rev. Earth Environ., 3, 197
Darroch S. A., Smith E. F., Laflamme M., Erwin D. H., 2018, Trends Ecol. Evol., 33, 653
Elliott-Graves A., 2023, Ecological Complexity. Cambridge University Press, Cambridge
Fauchez T. J et al., 2022, Planet. Sci. J., 3, 213
Ford Doolittle W., 2014, Biol. Philos., 29, 415
Gell-Mann M., 2002, Complexity and Industrial Clusters: Dynamics and Models in Theory and Practice. Springer, Heidelberg: Physica-Verlag HD, p. 13
Goldblatt C., 2016, preprint (arXiv:1603.00950)
Hawes I., Jungblut A., Matys E., Summons R., 2018, Geobiology, 16, 369
Hoffman P., 2016, Geobiology, 14, 531
Hull P., 2015, Curr. Biol., 25, R941
Hull P. M., Darroch S. A., Erwin D. H., 2015, Nature, 528, 345
Hyde W. T., Crowley T. J., Baum S. K., Peltier W. R., 2000, Nature, 405, 425
Jensen H. J., 2018, Eur. J. Phys., 40, 014005
Kauffman S., Levin S., 1987, J. Theor. Biol., 128, 11
Keppel G. et al., 2012, Glob. Ecol. Biogeogr., 21, 393
Kirchner J. W., 2002, Clim. Change, 52, 391
Kirchner J. W., 2003, Clim. Change, 58, 21

Laflamme M., Darroch S. A., Tweedt S. M., Peterson K. J., Erwin D. H., 2013, Gondwana Res., 23, 558

Landi P., Minoarivelo H. O., Brännström Å., Hui C., Dieckmann U., 2018, Population Ecol., 60, 319

Lechte M. A. et al., 2019, Proc. Natl. Acad., 116, 25478

Lenton T., Watson A., 2013, Revolutions that made the Earth. Oxford University Press, Oxford

Lenton T. M., Crouch M., Johnson M., Pires N., Dolan L., 2012, Nat. Geoscience, 5, 86

Lenton T. M., Dahl T. W., Daines S. J., Mills B. J., Ozaki K., Saltzman M. R., Porada P., 2016, Proc. Natl. Acad. Sci., 113, 9704

Lenton T. M., Daines S. J., Dyke J. G., Nicholson A. E., Wilkinson D. M., Williams H. T., 2018, Trends Ecol. Evol., 33, 633

Lovelock J. E., Margulis L., 1974, Tellus, 26, 2

Martin R. E., 1996, Palaios, 11, 209

McMahon S., Parnell J., 2018, J. Geol. Soc., 175, 716

Mills D. B., Boyle R. A., Daines S. J., Sperling E. A., Pisani D., Donoghue P. C., Lenton T. M., 2022, Nat. Ecol. Evol., 6, 520

Narbonne G. M., et al., 2012, The geologic time scale, 1, 413

Newman M., Sibani P., 1999, Proc. R. Soc. Ser. B, 266, 1593

Nicholson A. E., Mayne N. J., 2023, MNRAS, 521, 5139

Nicholson A. E., Wilkinson D. M., Williams H. T. P., Lenton T. M., 2018, MNRAS, 477, 727

Nicholson A. E., Daines S. J., Mayne N. J., Eager-Nash J. K., Lenton T. M., Kohary K., 2022, MNRAS, 517, 222

Payne J. L., Bachan A., Heim N. A., Hull P. M., Knope M. L., 2020, Interface Focus, 10, 20190106

Peltier W., Tarasov L., Vettoretti G., Solheim L., 2004, Was. DC Am. Geophys. Union Geophys. Monogr. Ser., 146, 107

Piovani D., Grujić J., Jensen H. J., 2016, J. Phys. A: Math. Theor., 49, 295102

Porada P. et al., 2016, Nat. Commun., 7, 12113

Quanz S. P. et al., 2022, A&A, 664, A21

Raup D. M., Sepkoski J. J., Jr, 1982, Science, 215, 1501

Ritort F., 1995, Phys. Rev. Lett., 75, 1190

Roach T. N., Nulton J., Sibani P., Rohwer F., Salamon P., 2017, Entropy, 19, 192

Rohde R. A., Muller R. A., 2005, Nature, 434, 208

Rosenfeld J. S., 2002, Oikos, 98, 156

Sahoo S. K. et al., 2012, Nature, 489, 546

Schrag D., Hoffman P., 2001, Nature, 409, 306

Scotese C. R., Song H., Mills B. J., van der Meer D. G., 2021, Earth Sci. Rev., 215, 103503

Seager S., Bains W., Hu R., 2013, ApJ, 775, 104

Solé R. V., Montoya J. M., Erwin D. H., 2002, Phil. Trans. R. Soc. Ser. B, 357, 697

Solé R. V., Saldaña J., Montoya J. M., Erwin D. H., 2010, J. Theor. Biol., 267, 193

Song H. et al., 2023, Nat. Commun., 14, 1564

Spaak J. W. et al., 2017, Ecol. Lett., 20, 1315

Stewart J. R., Lister A. M., 2001, Trends Ecol. Evol., 16, 608

Stewart J. R., Lister A. M., Barnes I., Dalén L., 2010, Proc. R. Soc. B, 277, 661

Wagner P. J., Kosnik M. A., Lidgard S., 2006, Science, 314, 1289

Ward P., 2009, The Medea Hypothesis: Is Life on Earth Ultimately Self-destructive? Princeton Univ. Press, Princeton, NJ

Wilkinson D. M., 2023, The Fundamental Processes in Ecology: Life and the Earth System. Oxford Univ. Press, Oxford

This paper has been typeset from a TEX/LATEX file prepared by the author.