Evolutionary Transitions in Individuality by Endogenization of Scaffolded Properties

Pierrick Bourrat

Abstract

The hierarchy of life is the result of a succession of evolutionary transitions in individuality (ETIs). During an ETI, individuals at a particular level of organization interact in such a way as to produce larger-level entities that become individuals in their own right. These new individuals are defined by their capacity to exhibit Darwinian properties of variation, differences in fitness, and heredity. One difficulty in accounting for ETIs is articulating how these properties are acquired at a higher level from the lower ones. Collaborators and I recently proposed the 'ecological scaffolding' model in which imposing an ecological scaffold (that is, a structure in the environment) on lower-level entities initiates an ETI. Here, I present a new model that extends this work. Within this new model, I propose a mechanism of scaffold endogenization, demonstrating that collectives can become resilient to the ecological scaffold being removed. This type of resilience is not observed in the ecological scaffolding model. However, classically, a biological individual would be regarded as an entity capable of withstanding environmental changes. Thus, the new model proposed here represents a step towards a more complete explanation for ETIs.

1 Introduction

Life is hierarchically organized. We are made of cells, which are made of organelles, some of which contain pieces of DNA. It is commonly accepted that life did not begin with fully fledged multicellular organisms or even unicells. Rather, the latter are the result of a number of transitions in which the free-living ancestors of lower-level entities initiated a class of interactions that led to the formation of higher-level entities (Maynard Smith and Szathmary [1995]; Michod [1999a]; Jablonka and Lamb [2006]; Bourke [2011]; Calcott and Sterelny [2011]; Clarke [2014]; van Gestel and Tarnita [2017]; Black et al. [2020]; Bourrat [2019], [2021b]). These newly formed entities acquired features that make them behave as cohesive wholes, which we then recognize as new individuals able to participate in higher-level processes of selection. Bourke ([2011], pp. 11–13) p proposed that over the course of the history of life on Earth, six types of evolutionary transitions in individuality (ETIs) have occurred. While ETIs will be discussed abstractly here, I will exemplify them with the fourth of Bourke's taxonomy, namely, the transition from unicells to multicellular organisms, which occurred in at least twenty-five different taxa (Grosberg and Strathmann [2007]).

Despite the fundamental importance of ETIs for the evolution of complexity, the mechanisms by which they occur are poorly understood. One influential conceptual model proposes that a change in the nature of fitness lies at the heart of ETIs. A prime example of this way of thinking can be found in the work of Michod and collaborators, who proposed the fitness transfer of ETIs (for a review, see Michod [2005]). In this model, an ETI has occurred when the fitness of lower-level entities (particles) has been transferred to the higher level (collectives). The same type of explanation can be found in Okasha's model of ETIs s (Okasha [2005], [2006], Chapter 8), some elements of which are inspired by the work of Michod and collaborators. This model proposes that the fitness of a collective at the outset of an ETI is defined as the average fitness of its constituent particles, and that this is no longer the case at the end of the transition. Following the distinction drawn by Damuth and Heisler ([1988]), Okasha proposed that this represents a transition from a multilevel selection 1 (MLS1) to a multilevel selection 2 (MLS2) process. More recently, without explicitly using the MLS1/MLS2 distinction, Okasha ([2016]) argued that collective fitness is 'derivative' at the beginning of an ETI but becomes 'primary' at its end. Michod and Okasha's model of ETIs represents what I will call an 'internalist' approach to ETIs where what does the explaining is a change in the internal properties of the entities undergoing the ETIs (in this case, their fitness).

While there is no doubt that once an ETI has occurred, the fitness of the constituent particles of a collective are aligned (Clarke [2014]; Clarke [2016b]; Bourrat and Griffiths [2018]), it is much less clear that changes in the nature of fitness do the explaining of ETIs for reasons discussed in Bourrat ([2015a], [2015b]) and Bourrat et al. ([unpublished]). At the very least, internalism represents only one possible mode of explanation for ETIs. According to another mode of explanation, which I call 'externalist', what does the explaining are changes in the ecology of a population. In this article, I present a specific model of this second mode of explanation for ETIs proposed by myself and collaborators (see Black et al. [2020]), expand it, and show some of its distinctive advantages over the internalist mode. I focus on a particular aspect of ETIs, which I call 'endogenization of scaffolded properties'.[1]

This article will run as follows: I start by presenting a set of criteria to assess whether an ETI has occurred. Following a number of authors, I argue that ETIs occur when units at a new level of organization exhibit the Darwinian properties of (1) phenotypic variation that (2) leads to differences in fitness and (3) that is heritable. Next, I present in more detail what distinguishes the two modes of explanation (internalism and externalism) for ETIs. I then present the externalist 'ecological scaffolding' model we proposed in Black et al. ([2020]). This model shows that starting from a population of particles organized in patches, with limited dispersion between patches (an ecological scaffold), imposing a selection regime at a timescale longer than the timescale at which the particles reproduce can lead to the formation of collectives that exhibit collective-like properties. As we recognized in Black et al., this model does not represent a full model of ETIs because, if the selection regime were to change, the collective-like properties (for example, a reduced growth rate) would disappear. From there, I present a new model that builds on Black et al.'s model by going one step further. I hypothesize that a trade-

---

[1] The word 'endogenization' is used here in a related but slightly different way to (Okasha [2006], [2021]; Bourrat [2014]). In those works, 'endogenization' refers to the strategy of explaining evolutionary phenomena that were previously taken from granted. For instance, in the context of ETIs, the biological hierarchy, which has traditionally been taken for granted, becomes part of what must be explained by evolutionary theory. In the present work, 'endogenization' is a process rather than an explanatory strategy.

off between growth rate and dispersion can cause the formation of collective entities that retain their properties even once the ecological scaffold is lifted. To demonstrate this point, I build an agent-based model that has the same basic features as in Black et al. but is based on a different modelling approach. After recovering results consistent with that of Black et al., I further show that when dispersion is correlated with growth rate, even if the ecological scaffold is lifted, the collectives that evolved due to the ecological scaffold can retain their scaffolded properties and remain insensitive to new (small) mutations indefinitely.

2 Characterizing Evolutionary Transitions in Individuality

An ETI involves a shift in individuality from the particle to the collective level. The notion of individuality is difficult in philosophy of biology and beyond (Pepper and Herron [2008]; Clarke [2010]; Guay and Pradeu [2015]; Sterner [2015]; Lidgard and Nyhart [2017]; DiFrisco [2019]). Therefore, I will not attempt to define this term rigorously here. However, within evolutionary theory, a 'Darwinian' individual is an entity able to participate in an evolutionary process 'in its own right' and, in particular, to participate in the process of evolution by natural selection (ENS) (Griesemer [2000]; Godfrey-Smith [2009]; Clarke [2013]; De Monte and Rainey [2014]). One prerequisite for this to occur is that this entity has properties that will confer to a population of its class the capacity to evolve by natural selection. Lewontin ([1970], [1985]), following a long tradition since Darwin (see Godfrey-Smith [2007]), recognized that for ENS to occur, a population should exhibit three conditions. These are (1) phenotypic variation that (2) leads to some differences in reproductive output or growth rate (fitness) between the entities exhibiting the variation and (3) that is passed on to offspring. Thus, following Lewontin's three conditions, to be a Darwinian individual is to have the capacity to exhibit (1) phenotypic variation, (2) differential fitness, and (3) heredity. Consequently, an ETI is an evolutionary process during which lower-level entities (particles) form higher-level entities (collectives) that acquire these three conditions.

However, this reasoning, despite its simplicity, is problematic. Typically, if the three conditions are exhibited at the particle level, they will also be exhibited by any set that contains more than one particle, whether these sets represent meaningful collectives in a biological sense or not. In other words, one can decide to group particles in any arbitrary sets, and these sets will exhibit Lewontin's three conditions (Glymour [2017]; Bourrat [2021a]). Thus, without supplementing the three conditions with some other criteria, the conditions are insufficient to distinguish real individuals or units of selection from pseudo or arbitrary ones. Consequently, they cannot be used to establish whether an ETI has occurred.

One way this problem manifests can be illustrated with what Okasha ([2006]) calls the problem of 'cross-level by-products'. To say that selection at the collective level is a by-product of selection at the particle level implies that the evolutionary change observed at the collective level should, in effect, be attributed to the level of the particle.[2] In other words, in situations of collective-level by-products, the collective level does not exist other than via arbitrary construction by an observer. One famous instance of collective-level by-product comes from Williams ([1966], p. 18), who told us that a herd of fleet deer is also a fleet herd of deer. Yet, herds are classically not regarded as individuals (that is, collectives that exhibit a sufficient degree of biological organization), despite satisfying Lewontin's three conditions broadly. A herd can exhibit different levels of fleetness (Condition 1), which can lead to being more or less successful in escaping predators and, thus, producing more or less offspring (Condition 2). Provided one can trace parent–offspring relationships between herds, one would find that a herd's level of fleetness is transmitted to its offspring (with more or less fidelity) (Condition 3).[3] There are many other examples of cross-level by-products. For instance, Sober ([1984]) used an example of collectives comprising individuals of different heights, in which height determines fitness. Assuming the height of an individual does not influence the height or fitness of other individuals in the collective, the average height of a collective and selection for height at the collective level are cross-level by-products of the particle level.

To remedy the failure of Lewontin's three conditions to discriminate arbitrary from genuine collectives, one can supplement the three conditions with the following condition: a genuine individual is generated by internal mechanisms as opposed to external mechanisms, or choices made by an observer (no mechanisms).[4] In other words, what distinguishes a biological indi-

---

[2] This problem is also faced by the Price equation. Surprisingly, it is much less often mentioned for Lewontin's three conditions, despite the near equivalence between the two (see Okasha [2006], Chapter 1).

[3] The difficulty of accomplishing this in practice does not disqualify the conceptual point.

[4] For another compatible solution to this problem, see (Bourrat [2021a], [2021b]).

Properties	Definition
Pseudo-Darwinian	The entity is attributed the capacity to exhibit phenotypic variation, fitness difference, and heredity only due to an observer's choice (that is, without it being the result of individuating mechanisms).
Darwinian-like	The entity is attributed the capacity to exhibit phenotypic variation, fitness difference, and heredity due to external (that is, ecological) individuating mechanisms.
Darwinian	The entity is attributed the capacity to exhibit phenotypic variation, fitness difference, and heredity due to internal individuating mechanisms that provide the entity some resilience to perturbations.

Table 1. Definition of the different types of properties related to Lewontin's three conditions an entity can exhibit

vidual from an arbitrary one is the existence of some internal mechanism(s) or properties of this entity that generate the three conditions. One implication of this additional condition is that an individual is resilient to some external perturbations due to its internal mechanisms being isolated from the environment. Further, I argue that the broader the range of perturbations an entity can sustain, the higher its score on individuality. To be clear, I will call Lewontin's conditions 'Darwinian properties' when a collective exhibits them due to internal mechanisms. Corresponding to Darwinian properties are 'Darwinian individuals'. Following our proposal in (Black et al. [2020]), when the three conditions are exhibited by a collective for reasons that are purely exogenous to it, I will call them 'Darwinian-like properties', and the collective a 'Darwinian-like' individual. When the conditions are exhibited for purely arbitrary reasons (such as choices made by an observer), I will refer to 'pseudo-Darwinian' conditions and 'pseudo-Darwinian individuals'. Thus, I submit that an ETI occurs when initially free-living particles in a population interact in such a way that they produce collectives exhibiting Darwinian properties, which did not exhibit them to begin with (see Table 1).

There are some clear parallels (and some inspiration on my part) between my proposal and that of Clarke ([2013], [2014], [2016a]). According to Clarke, once an ETI from particles to collectives is complete, the collectives exhibit heritable differences in fitness. I interpret this to mean that they exhibit Lewontin's three conditions,[5] and there are some individuating mechanisms at that level. More precisely, for Clarke, individuating mechanisms are divided into

---

[5] Note that, as highlighted by Okasha ([2006], p. 13), heritable variation should not necessarily be variation in fitness but rather in phenotypes.

demarcation mechanisms (which maintain variation among the collectives of the population) and policing mechanisms (which retain the collective's integrity). Clarke tells us that to be a unit of selection or a Darwinian individual, a collective must be able to exhibit Lewontin's three conditions; further, they do this due to a combination of demarcation and policing mechanisms. Thus, being an individual results from some objective processes in the world (individuating mechanisms), not some arbitrary decision made by an observer. If there is no objective mechanism at the herd level leading a herd of deer to generate heritable phenotypic variation and differential fitness, it should not count as an individual. Clarke's proposal provides a simple yet powerful way to detect when an entity ought to be regarded as a Darwinian individual.

Clarke's strategy, like mine, is part of a long tradition in philosophy of biology that Griesemer ([2005]) has called 'generalization by abstraction'. Griesemer provides one important criticism of this strategy—a purely functionalist account such as Clarke's does not provide a guide for choosing the material properties that will permit one to assess whether an entity is an individual. Yet, this step is critical for obtaining an empirically adequate account. Both Sterner ([2015]) and DiFrisco ([2019]) have developed a version of this criticism in their discussions of Clarke's account. Further, DiFrisco also provides a tentative solution. He proposes that Clarke's through-and-though functionalism is grounded in an inclusive disjunction of material realizers. Each realizer (concrete biological mechanism) can perform the function of being an individuating mechanism. However, he stresses that 'function' here ought to be defined in terms of capacity rather than actual performance (its actual effects). In other words, traits can malfunction and, nevertheless, be ascribed a function, which is not the case if function is defined as actual effects. One account of function that allows for such ascription to poorly performing traits is the selected-effect account of function (for an introduction and discussion of this account, see Garson [2019]). This solution renders Clarke's approach (and, consequently, mine) viable.

An essential difference to note between Clarke's account and mine is that for Clarke, any mechanism, whether external or internal to the collective, can be used to assess whether a collective is an individual. This is clear when she writes: 'I remain agnostic, at this stage, about whether individuals must possess individuating mechanisms intrinsically—within their own skins, so to speak. Perhaps it is sufficient, especially in the early stages of a transition, for the mechanisms to exist in the environment, so long as they are stable enough that their effect is heritable' (Clarke [2013], p. 427, Footnote 53). Elsewhere, she proposes that mechanisms leading to aggregation, assortment, or (again) fitness correlations between particles are all individuating mechanisms without specifying their causal origin (Clarke [2014]). This implies that a situation in which a collective that satisfies Lewontin's three conditions for purely exogenous reasons could count as Darwinian individuals having undergone an ETI for Clarke—whereas, for me, they are only Darwinian-like individuals.

Clarke's motivation is clear—she wishes to avoid having to explain ETIs by invoking higher-level properties as causes because they are the consequence of ETIs and, thus, in need of explanation rather than part of the explanation. This leads to what she calls a 'chicken and egg problem' (see also Griesemer [2000]; Veit [2021]). I agree with her that exogenous mechanisms can be the initiator of ETIs. However, I consider that the mechanisms producing the three conditions have to be internalized or endogenized for an ETI to be genuine. Conversely, Clarke does not require this step.

There are two reasons why I regard the endogenization step as necessary for achieving an ETI. First, if the three conditions are exhibited solely due to external mechanisms, the cross-level by-product problem might not be solved. Again, suppose there is a population of deer in the savanna, with scattered patches of grass, and that a single herd of deer grazes a single patch at a time. We assume that the sole reason the deer are organized in herds is the presence of patches of food. Occasionally, a predator attacks a herd. Fleeter herds are more successful and have more offspring (whether counted by the number of deer or herds) with similar traits as their parent(s). This situation fits Clarke's model. There is an individuating mechanism for the herd and, thus, one would have to conclude that the herd is an individual, although biological wisdom would tell us that it is, at best, a Darwinian-like individual following my terminology.[6]

Second, without internalized individuating mechanisms, any change (even small) in the

---

[6] Note that there is no policing mechanism here; thus, the herds are, strictly speaking, not individuals following Clarke's model, which requires such a mechanism. However, this should not distract us from the main point. One could, indeed, imagine other scenarios, such as those presented in the following section, in which a single particle founds a collective. The process by which a bottleneck is created represents a policing mechanism for everything else being equal, leading to less competition within a collective. The bottleneck can be purely due to ecological conditions.

factors of the environment might lead the population to return to a situation where there are no collectives anymore. As mentioned earlier, I consider that to be an individual is to have some capacity of resilience to environmental changes.[7] These differences to one side, and using DiFrisco's supplementation, I am broadly in agreement with Clarke's analysis.

3 Internalism and Externalism Regarding ETIs

In the previous section, I argued that for an ETI to occur, an initial population of free-living particles should ultimately reach a stage at which internal individuating mechanisms lead to the maintenance of collectives exhibiting heritable phenotypic differences that produce differences in fitness. In short, collective-level entities should acquire Darwinian properties.

One influential model of ETIs, which seems to satisfy this process, has been proposed by Michod and collaborators (for a review, see Michod [2005]). This model proposes that Darwinian properties are acquired by a transfer of fitness from the particle to the collective level during the ETI. In the transferral process, the lower-level entities relinquish their individuality. At the end of the fitness transfer, collectives exhibit Darwinian properties but the particles do not because they are no longer fitness bearers. As mentioned in the introduction, the model proposed by Okasha ([2006], Chapter 8) based on the MLS1/MLS2 distinction follows a similar line of reasoning. Michod and collaborators propose two mechanisms by which fitness can be transferred. One involves the coupling of two loci (locus modifier model; see Michod and Roze [1997]; Michod [1999b]), the other relies on life history theory (see Michod et al. [2006]; Doulcier et al. [unpublished]).

The fitness transfer model and other models inspired by it are instances of internalist explanations for ETIs. I define an internalist explanation, roughly following Godfrey-Smith's ([1996]) terminology, as an explanation in which internal properties of the entities exhibiting the trait to be explained do the explaining. For an externalist explanation, the external properties do the explaining. The distinction between internalism and externalism is relative. When dealing with biological systems, any trait is the result of an interaction between the environment and internal

---

[7] To be fair, Clarke ([2013]) mentions that for environmental mechanisms to be individuating mechanisms, they would have to be stable enough to be heritable. However, it is hard to see how external stability could be achieved without its coupling with some internal properties that would control the external mechanism. However, this would *de facto* render the mechanism (at least partly) internal.

properties. However, the properties that 'make a difference' for this phenomenon can be either internal or external to the entities exhibiting the trait to be explained. Thus, to be an internalist about ETIs is to believe that the mechanism used to explain the transitions relies ultimately on a change in the internal properties of the entities undergoing the transition.

We saw in the previous section that to be a Darwinian individual is to have some internal mechanisms leading to the production of Darwinian properties. Yet, the problem is to explain them away without requiring that they are initially present (Clarke's chicken and egg problem). The fitness transfer model solves this conundrum by proposing a change in the particles' internal properties—namely, their fitness—with no or little mention of changes in the environment. In principle, there is nothing wrong with the internalist mode of explanation. Mutations, which are internal changes, once combined with selection, could lead to ETIs in a single selective environment. Further, internalist models about ETIs provide a solution to the acquisition of Darwinian properties at the collective level. By being transferred from one level to the other, fitness becomes a heritable internal collective property that is resilient to perturbations.

Despite the appeal of the fitness transfer model, there are some problems with it. First, fitness is a difficult and debated concept (for a survey, see Abrams [2012]). Consequently, this makes the notion of fitness transfer difficult to define precisely and consistently. Second, at the end of a transition, particles are assumed to have nil fitness, which is also difficult to make sense of at face value since the cells (or at least some of them) of a collective—say, a multicellular organism—reproduce at the end of a transition (Godfrey-Smith [2011]; Bourrat [2015a]). Partly in response to this line of criticism, Shelton and Michod ([2014], [2020]) have argued that the fitness transfer model ought to be understood as a distinction between the fitness that a cell would have in isolation and in the group. However, understood as such, fitness is only counterfactually transferred, which is problematic since the notion of transfer implies that it actually happens. While the counterfactual approach provides an explanation of why cells become part of a multicellular organism, there is no literal transfer of fitness from one level to the other (Doulcier et al. [unpublished]; Bourrat et al. [unpublished]). Finally, in some versions of the transfer fitness model, the authors have been accused of presupposing the individuating mechanisms at the collective level rather than them being the outcome (Clarke's chicken and egg problem). These difficulties will not be discussed further here. However, let me simply say that they provide sufficient motivation to find alternative explanations of ETIs.

Recently, my collaborators and I proposed an externalist model for ETIs (Black et al. [2020]). Rather than assuming that something changes fundamentally in the internal properties of the entities undergoing an ETI (as in the fitness transfer model), it assumes that an ETI can be driven by the ecology of those entities. To be clear, internalist models need not oppose externalist ones. It is plausible that ETIs involve both internal and external changes (see Section 7). However, the externalist approach shifts the perspective on ETIs by proposing that what initiates them comes from a change in ecology. It permits us to explore a different possibility within the space of possibilities for ETIs. Insofar as ETIs have occurred across multiple levels of biological organization and among many different taxa, it is plausible that the transitions occurred by different mechanisms; thus, an externalist model is welcome even if it only explains a subset of all ETIs.

4 Black et al.'s Model

The model my collaborators and I proposed in (Black et al. [2020]) assumes a population of cells growing at different rates in a structured environment made of patches. Patches contain some resources. If the resources are depleted before a particular time T, which corresponds to many cycles of cell replication, the cells on the patch die. At the end of a patch generation (T) (that is, a discrete generation), cells can found a new patch from a single cell, as a function of the number of cells they had at the end of the previous patch generation (the higher number of cells at the end of T, the greater the opportunity to found new patches). Crucially, this mode of reproduction is the only way for a cell to invade a new patch. In other words, there is no migration between patches except over the timescale T.

A classical Darwinian picture would want that when cells are in competition for resources, the faster growing cells invade the population. However, this is not what is observed. If T is sufficiently long, the cells that grow more slowly can invade the population. This is so because fast-growing cells, more often than not, over-exploit their patch resources before having the opportunity to found a new patch; thus, they are selected against in the long term. Further, the fact that the population is highly structured makes any short-term benefit 'local' both in time and space with no possibility to propagate to the wider population. One consequence is that, in the long term, the slow-growing cells appear to be the fast-growing ones. When we refer to 'slow' and 'fast' growing cells, this implicitly assumes a reference to the short-term or well-mixed environment.

This model shows that if cells are organized in collectives on a patch that are imposed by the environment (that is, no migration between patches, selection of patches over a long timescale, and a bottleneck between each event of patch founding), which we term a 'scaffold', it is possible to evolve a phenotype that would otherwise be selected against.[8] One interesting aspect of our model is that it makes the obtention of Lewontin's three conditions at the collective level entirely non-mysterious—they result from externally imposed conditions on the population.

The model we presented in (Black et al. [2020]) could be sufficient for Clarke to consider the collectives living on a patch as individuals. However, these collectives are not individuals following my requirements. Patches or, more accurately, the collective of cells living on a patch, do indeed exhibit Lewontin's conditions at that level. There can be variation in the average growth rate of the collectives, which is transmitted (with small mutations) from one patch generation to the other and can lead to differences in the number of collective offspring. However, the mechanisms producing Lewontin's three conditions at that level—that is, the bottleneck between events of patch foundation (policing mechanism) and a strict boundary between patches preventing between-patch migration (demarcation mechanism)—are purely external to the collective. The properties are Darwinian-like, as we characterize them—they are not Darwinian. Thus, it does not represent a full model of an ETI where the individuating mechanisms have been endogenized; rather, it can be regarded as a way of initiating the ETI.

In the previous section, I identified two reasons why a purely exogenous individuating mechanism will be insufficient to explain away an ETI. First, it might not solve the cross-level by-product problem. This problem is not present in the model we presented in (Black et al. [2020]). A cross-level by-product occurs when there is no interaction between the particles in the production of a collective phenotype. However, in our model, the long-term growth rate

---

[8] For more on the concept of 'scaffold' in evolution, see (Caporael et al. [2013]; Veit [2021]).

of a cell is different from what it would be if it were living freely (that is, with no interaction with other cells). The cells do interact with one another, but this interaction is mediated by the environment (that is, the resources of the patch). Thus, a collective here is not analogous to a herd of deer where there would be no interaction between the deer. However, I also argued that to be considered a Darwinian individual, an entity should be able to withstand some ecological changes. In the model presented in (Black et al. [2020]), if the scaffold was to be lifted, the collectives would quickly disappear. This is because cell migration between patches would be allowed over the cell replication timescale. As a result, a mutation increasing the growth rate would mean a higher chance of invading a new patch. The population would soon be invaded by fast-growing cells, with the possibility of ultimately collapsing due to global over-exploitation of patches. At any rate, the collective would disintegrate, and any way of grouping cells with one another would become arbitrary. In other words, collective-level properties could only be of the pseudo-Darwinian kind (see Table 1).

5 The Endogenized Scaffold Model

The model presented in (Black et al. [2020]) permits us to understand how an ETI can be initiated by exporting individuating mechanisms to the environment. However, it does not represent a full model of ETIs because the collectives have not endogenized those mechanisms. In this section, I propose an agent-based model that extends the original ecological scaffolding model. A key result of this model is that, under some conditions, collectives are able to evolve properties that makes them resilient to the removal of the scaffold. This occurs through the endogenization of an individuating mechanism. Thus, this model presents a mechanism by which Darwinian properties at the collective level are able to evolve from non-Darwinian collectives through Darwinian-like collective level properties. To be upfront, we were clearly aware in (Black et al. [2020]) that our model is insufficient to explain ETIs fully and that only by removing the scaffold can an ETI be deemed complete. We proposed a number of potential avenues by which this could be initiated without modelling them, such as the evolution of a storing resource via seeds or the evolution of developmental processes (for an exploration of this idea, see Doulcier et al. [2020]). While these are clearly worthwhile paths to explore, I propose here another avenue to endogenization, which relies instead on pleiotropy.

The new model assumes a grid of patches with some resources and on which cells live. These resources are limited and perishable; they are limited because a patch can only sustain a certain number of cells. If the patch is over-exploited or if the resources have perished, all the cells of the patch die. If a patch contains no cells, it is replenished with resources. Cells reproduce perfectly except for some random mutations, which can occur during replication and cause an increase or decrease of a cell's growth rate by 1% of a maximum growth rate, the latter of which is set arbitrarily. It is further assumed that a cell's growth rate depends purely on its genotype. Once a cell has replicated, it can either stay on the patch or move to another patch. Whether it stays or moves depends on the value of a second trait, which is a pleiotropic effect of the genotype. If the growth rate of cell reproduction is high, the probability of leaving the patch and colonizing a neighboUr patch is high. Conversely, if the growth rate is low, the probability of leaving the patch is also low. The model assumes that whether a cell is able to colonize the neighbour patch also depends on whether this patch is free and, if it is not, whether the cells on it score high on the pleiotropic effect. If they score high, the probability of colonizing is low (but non-nil); if they score low, the probability is higher. One realizer of the relationship between growth rate, probability to leave the patch, and ability to resist invasion could be cells producing glue. Producing glue is costly and, thus, lowers the growth rate and prevents the cells from leaving the patch, following the 'wrinkly spreader' (*Pseudomonas fluorescens*) model proposed by Rainey and collaborators (Rainey and Rainey [2003]; Rainey and Kerr [2010]; Hammerschmidt et al. [2014]; Rose et al. [2019]). We could also hypothesize that the glue prevents new migrant cells from colonizing a patch even though this is not part of the wrinkly spreader model.

Finally, we assume that at each T generation, the cell replication cycle is different due to some periodic environmental changes. In those generations, cells have no opportunity to replicate and move to another patch except for one cell per patch chosen randomly among the cells with a level of stickiness above 1.5 or, equivalently, a growth-rate below 0.5 (assuming this patch contains such cells), and only if the neighbour patch is free or if not, if the mean value for stickiness is sufficiently low (which is determined probabilistically). This second mode of reproduction could be regarded as a form of specialization of some cells, which is akin to resulting from a form of collective-level proto-development, allowing one of the cells to become a germ cell of the collective and to found a new collective by producing a 'propagule'. A NetLogo version (Wilensky [1999]) of the model is provided as supplementary material.

With this setting in place, we assume three regimes. In the three regimes, the model is run for 30,000 cell generations. In the first regime, the 'no-scaffold' regime, nothing is different from what is described above. This can be regarded as a control. In the second regime, the 'scaffold' regime, migration between patches is only allowed when there is the opportunity for collective to produce a propagule (that is, at generations that are a multiple of T). In principle, cells could migrate between patches when they reproduce, as observed under the no-scaffold regime. Nevertheless, this is prevented by the ecological scaffold. This regime represents a situation similar to the model we proposed in (Black et al. [2020]). However, there is one difference: the patch generations are not discrete. Finally, the third regime, which I call the 'reverted' regime, is a mix of the two others. First, the cells are under a scaffold regime for 1500 cell generations. After that time, the scaffold is lifted, and the regime becomes a no-scaffold regime until the end of the run.

Under each regime, the timescale T is varied from $T = 3$ to $T = 10$, and the maximum collective size before patch resources are over-exploited from $S = 4$ to $S = 20$ with increments of two cells per condition. This leads to 72 conditions within each regime. Each condition is repeated 01 times, and the average of all the repeats that completed the 30,000 cell generations is taken. For each run, there is initially a single cell per patch, and the overall number of patches is 200 (a torus of twenty by ten patches). Initially, each cell is attributed a random float genotype taken from a normal distribution with a mean of 1 and a standard deviation of 0.25. If the genotype has a value below 0 or above 2, which are arbitrary boundaries for the purpose of the model, they are set to 0 and 2. If a cell has a genotype inferior to 0.5, it exhibits the capacity to produce a propagule over T. The age of patches is also randomized as an integer between 1 and 30. Resources perish when a patch reaches the age of 30 cell generations. An illustration of the model is presented in Figure 1 and further details are given in the info tab of the NetLogo model provided as supplementary material.

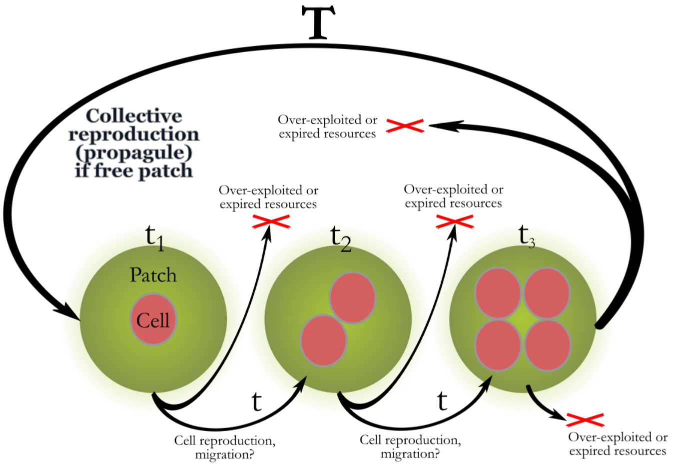

Figure 1. Illustration of a single collective generation (T) assuming three cell generations (t_1, t_2, and t_3) for the endogenized scaffold model. t represents the timescale for a cell division; 'migration?' indicates that the possibility of migration between patches exists under the 'no scaffold' regime but not under the 'scaffold' one. A free neighbour patch is defined either as a patch on which there is no cell, or a patch on which the growth rate of its cells have a high growth rate and consequently a low level of stickiness.

6 Results

6.1 Evolution without scaffold

Running the model under the no-scaffold regime leads invariably to the outcome observed in Figure 2 (blue line). The average genotype (which tracks cell growth rate) quickly increases until it reaches the maximum level (close to 2) and stays there until the end of the run. This result is observed at every timescale (from $T = 3$ to $T = 10$) and for every maximum collective size (from $S = 4$ to $S = 20$). Figure 3a displays the average genotype over the last 500 cell generations of a mean run for each timescale and maximum collective size before the patch resources are over-exploited under the no-scaffold regime. As can be observed, all growth rates (tracked by the genotype) are close to 2.

These results show that in the absence of a scaffold, the cells with a maximum growth rate (and a lower degree of stickiness and incapacity to produce propagules) are favoured by natural selection, as one would expect. Thus, no Darwinian or Darwinian-like properties at the collective level are observed.

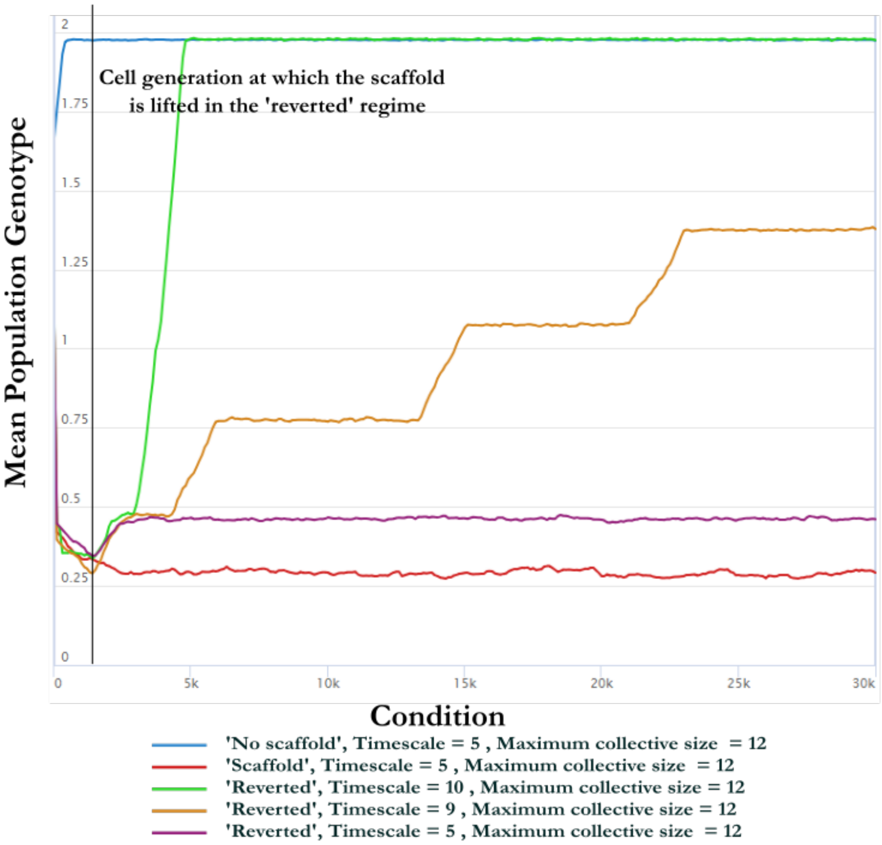

Figure 2. Comparison of different evolutionary trajectories in different regimes of the endogenized scaffold model. Under the 'no-scaffold' regime (blue line), the average genotype, which perfectly tracks cell growth rate, quickly increases near the maximum value, indicating that the cells are free-living and do not exhibit any phenotype as part of a collective. Under the 'scaffold' regime (red line), the average cell growth rate quickly decreases and remains as such indefinitely, indicating that the cells behave as part of a collective. Under the 'reverted' regime, when the timescale and maximum collective size before patch resource over-exploitation are low (purple line), the growth rate remains below 0.5 and, thus, similar to what is observed under the scaffold regime. However, if the timescale (or maximum collective size, which is not represented in the figure) is larger (orange and green lines), the collective can revert to a population of cells with a maximum growth rate. The larger the timescale and maximum collective size, the quicker this reversion occurs.

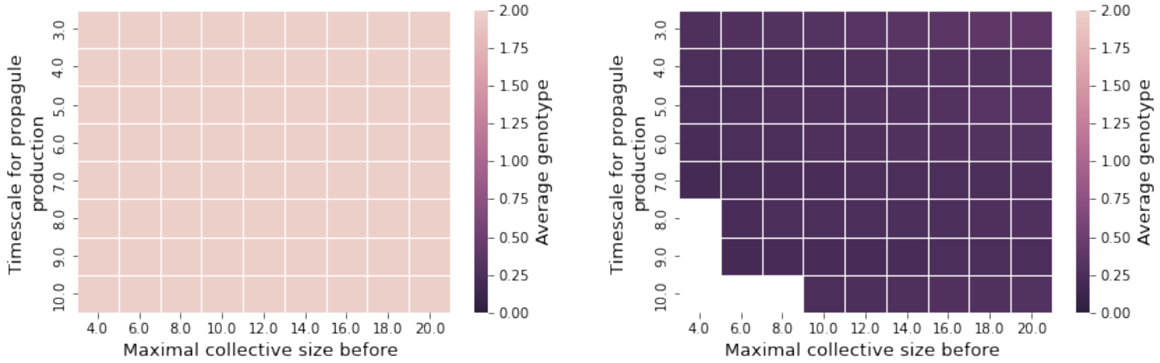

(a) No Scaffold **(b)** Scaffold

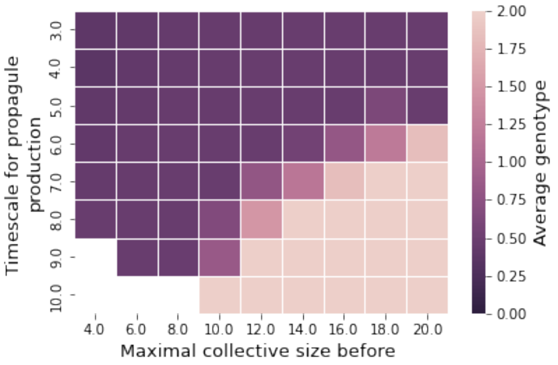

(c) Revert

Figure 3. Heatmaps showing the average genotype (which tracks growth rate, stickiness, and capacity to produce propagules over timescale T) over the 500 last cell generations out of $30,000$ with 'maximum collective size' before patch resource over-exploitation varied from four to twenty cells (with increments of two cells) on the x-axis and 'timescale for propagule production' varied from $T = 3$ to $T = 10$ on the y-axis. (a), (b), and (c) correspond to the 'no scaffold', 'scaffold', and 'reverted' (that is, scaffold and then no scaffold) conditions, respectively. White cells indicate that the set of conditions are not viable (the population always collapses before the end of the run).

6.2 Evolution with scaffold

Under the scaffold regime, the result is predominantly invariable; however, things look very different compared to the no-scaffold regime. The cells, with an initial average growth rate of 1, quickly reach a low growth rate in most conditions and remain at this growth rate throughout the run (below 0.3), as can be observed in Figure 2 (red line). This result is also observed at most timescales (from $T = 3$ to $T = 10$) and for every maximum collective size (from $S = 4$ to $S = 20$). Figure 3b displays the average genotype over the last 500 cell generations of a mean run for each timescale and maximum collective size. As can be observed, most growth rates (tracked by genotype) for the scaffold regime are lower than 0.3. Conditions where the cell colour of the heatmap is white indicates that these conditions are not viable. This occurs when the timescale for propagule production is over eight cell generations and the maximum collective size before patch resource over-exploitation is smaller than ten cells. It can be explained as follows: If the timescales over which propagule production can occur is longer than the time it takes for the collective to reach the maximum size and over-exploit the resources, the cells of the collective die before they can produce a successful propagule. The smaller the maximum collective size, the higher the likelihood this will occur.

These results indicate that the imposition of an ecological scaffold in which cells cannot invade another patch, except every T generation if they have a certain size, selects for sticky cells that can be regarded as initiating an ETI. This is broadly consistent with the results we obtained in (Black et al. [2020])—the imposition of a scaffold leads to Darwinian-like entities.

6.3 Removing the scaffold

Finally, moving on to the reverted regime (with which I am mostly concerned here), we observe that, in some conditions, once the scaffold is removed, cell growth recovers the value that it would have had without the scaffold (see Figure 2, green line). However, under other conditions, it remains at a low value despite the lack of a scaffold (see Figure 2, purple line). Finally, for the same reasons as those provided for the scaffold regime, some conditions are not viable. Figure 3c shows that whether the growth rate reverses to that observed under the no-scaffold regime outcomes depends on both the timescale T at which propagule production can occur and the maximum collective size before patch resource over-exploitation. If T is not too long and the maximum collective size is not too large, the cell growth rate does not increase past 0.5 after the scaffold is removed. These values are slightly higher than in the values obtained under the scaffold regime—but well below the values obtained under the no-scaffold regime—and are due to the fact that collectives can only produce a propagule if they contain cells with a growth rate below 0.5. In situations where the growth rate reverts to the level of that observed under the no-scaffold regime, it takes more time to reach this level when T is smaller and S is large. This point is illustrated by comparing the orange and green lines in Figure 2, which correspond to conditions under the reverted regime with $T = 9$ and $T = 10$, respectively (in both cases, $S = 12$ is twelve). As can be seen, it would take more than 30,000 generations for a condition where $T = 9$ and $S = 12$ (orange line) to recover the growth rate observed under the no-scaffold regime (blue line).

The results obtained under the reverted regime can be explained as follows: In some cases, when a population is composed of cells with a low growth rate, even in the absence of ecological constraints on between-patch migration (no scaffold), the long-term probability of having more descendants for those cells than competitors is higher by producing a propagule over T (which involves staying on the patch as part of the collective) than by reproducing individually over t. Of course, if the growth rate of the cells magically became higher (and, thus, less sticky), they would easily invade the population, as under the no-scaffold regime. However, reaching this phenotypic state is prevented by the cost of mutating—and as a result of over-exploiting the patch—that is imposed on those cells. In other cases, when T or S is large, the long-term benefit of reproducing as part of a collective via a propagule over T becomes lower than that of reproducing individually because the cost of over-exploitation becomes lower or the chance to produce as part of a collective becomes less frequent. From then on, the dynamics revert to the no-scaffold regime. The model and analysis provided here are preliminary; further work should confirm these results using analytical methods or more sophisticated simulations.

7 Discussion

The results obtained under the reverted regime highlight the significance of the interaction between the environment and the internal properties of the entities forming a population during ETIs. Starting from a setting where cells are interacting with the environment (patches) and have the potential to grow fast, they can evolve properties that make them behave as part of a collective. In some cases, the net long-term reward is higher for staying as part of a collective and having a chance to produce a propagule than the benefit of growing fast and invading other patches in the short term as single cells. The mutations leading to a higher growth rate, despite seemingly creating a short-term advantage, can nevertheless imply a lower growth rate in the long term. Once the scaffold is lifted, because dispersion and growth rate are coupled, similar constraints can nevertheless apply—while there are no external boundaries preventing the cells from moving to another patch beyond niche availability, they do not necessarily evolve the capacity to do so. This is because it would imply growing faster and, under some conditions, condemning the patch to be over-exploited before the cell or its descendants can invade another patch.

Using a proof of principle approach, my model begins with a population where collectives exhibit Darwinian-like properties, and shows that the possibility of endogenizing those to become genuine Darwinian properties (assuming a small number of extra assumptions) exists. To my knowledge, the transition from Darwinian-like to Darwinian properties has been underexplored in the literature. One exception comes from Libby et al. ([2016]). They present a model for ETIs where organisms can be in two different environments. One environment favours unicellularity while the other favours multicellularity. They show that the longer a population of organisms remains in the environment favouring multicellularity and can accumulate beneficial mutations for the collective in this environment, the longer it takes for these collectives to revert back to a unicellular phenotype once the collectives are then placed in the environment favouring unicellularity.[9] They call this process 'ratcheting' (see also Libby and Ratcliff [2014]). They then consider a second type of ratcheting where mutations alter the probability of switching from a unicellular to multicellular state. Finally, they combine

---

[9] In actuality, they only model cells, not collectives of cells, but this need not concern us here.

the two types of ratcheting to show that both types of ratcheting, once combined, can work synergistically.

Libby et al.'s ([2016]) model bears some resemblance to the model presented here in that it addresses the problem of reversion to unicellularity. However, I argue that it is less ecologically driven since population structure is temporal only in their model. Conversely, in my model, it is also spatial, which takes into consideration that collectives are spatially bounded entities. Second, their model does not demonstrate an irreversible transition to multicellularity within the scope of the parameters used. Switching between the multicellularity-favouring environment and the unicellularity-favouring environment can be regarded as similar to lifting the scaffold in my model. However, when this occurs in their model, the result (in both types of ratcheting scenario) is ultimately a switch back to unicellularity and, thus, a loss of Darwinian-like properties at the collective level. In contrast, there are indications (although this would require further investigation) that in my model, the Darwinian-like properties are transformed into Darwinian properties because the collective-level properties can be retained indefinitely in an otherwise unicellularity-favouring environment (no-scaffold regime) due to pleiotropic effects. To be clear, the goal here is not to criticize Libby et al.'s ([2016]) model, but rather to show that there potentially exists a multitude of roads towards the transition from Darwinian-like to Darwinian collective properties.

Following the experimental work by Hammerschmidt et al. ([2014]), I have proposed that the production of collective-level properties manifesting in a lower (short-term) growth rate and higher stickiness at the cell level could manifest biologically in the production of a glue that is costly to produce or an incomplete cell division, respectively. However, these two cell phenotypes should not be regarded as exhaustive for the production of Darwinian-like and Darwinian collective-level properties. For instance, in (Black et al. [2020]), we proposed that one way cell short-term growth rate is lowered is by the production of cells that do not have direct descendants (that is, effectively, somatic cells). In situations where resource over-exploitation leads to collective death before any opportunity to disperse to other patches, everything else being equal, the existence of such cells is beneficial in the long term. There is a myriad of mechanisms that could produce a reduced cell growth (up to the point where cells do not reproduce) associated with long-term benefits.

To give a further example, 'stickiness' in the model presented here could be replaced with 'defence of the patch' since a patch containing a high number of cells with a low growth rate is more difficult to invade than a patch where the growth rate is higher—recall that this is because the resident cells exhibit a pleiotropic phenotype that renders patch invasion more difficult. A classic, albeit much more complex, example of this phenomenon is the evolution of an immune system in some multicellular organisms, with some cells specialized in the recognition and elimination of entities that are sufficiently different from the resident cells.[10] Another example, during the transition from multicellularity to eusocial colonies, is the emergence of individuals that defend the colony against foreign intruders. Thus, we see that the abstractness of the endogenization model could, with slight modifications, accommodate a variety of evolutionary scenarios at different levels and be empirically tractable.

One further strength of my approach is that it connects with well-established evolutionary frameworks, particularly kin selection and niche construction. Classical models of kin selection often only assume intra-generational social effects from actors to recipients. In situations of limited dispersion, the inclusive fitness of an entity must also take into consideration the competition for resources between entities of same type living in close proximity—in a large, well-mixed population, this phenomenon does not exist (Bourke [2011], pp. 48–49). This can lead to situations where the benefits provided to relatives by altruism are outweighed by the resulting competition between relatives. Lehmann ([2007]; see also Lehmann [2008]) proposed a formal model where the addition of inter-generational (post-humous) effects can facilitate the evolution of altruistic traits in situations of limited dispersion. Post-humous effects are effects occurring after the death of an entity and, consequently, affecting its descendants. While Lehmann's model is framed within the kin selection framework and, thus, does not address the evolution of collective-level properties per se, it relies on the infinite island model proposed by Wright ([1931]). In this model, an infinite population of alleles is subdivided into groups of finite and equal sizes that can exchange at random a certain proportion of alleles at each generation. The endogenization of scaffolded properties model presented here, particularly under

---

[10] For a philosophical analysis of the relationship between immunity and biological individuality, see (Pradeu [2012]).

the no scaffold regime, although more mechanistic and ecological, shares a number of essential features with the infinite island model. In both the model we presented in (Black et al. [2020]) and the endogenization model presented here, the way a lower growth rate can evolve is through the existence of post-humous beneficial effects, which is compatible with the results from Lehmann's model. Indeed, in both models, the over-exploitation of resources on a patch, which can be selected against in the long run, results from the cumulative effects of cells consuming resources at a high rate over several generations.

The idea that post-humous effects can be selected for or selected against connects with the literature on niche construction, as emphasized by Lehmann. Niche construction theorists have proposed that niche construction is a process by which organisms—or, more generally, entities—in a generation modify their environment in such a way that it impacts the future success of their descendants (Lewontin [1983]; Odling-Smee et al. [2003]; Scott-Phillips et al. [2014]; Laland et al. [2016]; Tanaka et al. [2020]). Negative niche construction corresponds to effects that cause the maladaptation of future generations to their environment (for example, polluting the environment). In contrast, positive niche construction corresponds to effects that increase the evolutionary success of subsequent generations. Accordingly, resource over-exploitation can be conceptualized as a form of 'negative' niche construction, while a lower rate of exploitation can be regarded as a form of 'positive' niche construction.

While the exploitation of resources at different rates can be examined from the lens of niche construction, this is not where the main interest of this framework lies with respect to the endogenization of scaffolded properties. Rather, to understand the relevance of niche construction, consider a modified version of the endogenization model presented in Section 5. Suppose now that instead of the scaffold being purely exogeneously imposed, the cells produce a phenotype that becomes part of their environment,[11] persists over multiple cell generations, and creates the scaffold for the endogenization of collective level properties. We could imagine, for instance, that dead cells or waste products form a barrier surrounding a patch, limiting the growth of cells inside this patch and preventing the cells of other patches from entering. This case of negative niche construction (initially) could cause the emergence, over time, of selection pressures

---

[11] In the most up-to-date version of the gene selectionist framework (see Haig [2012]; Lu and Bourrat [2018]; Ågren [2021]), this part of the environment, from the perspective of a replicator, is simply an extended phenotype.

equivalent to those of the ecological scaffolding model, but initiated by the lower level entities themselves. Initially, the accumulation of this product would be detrimental to the cells living on the patch since it would, ultimately, condemn them. However, it could also prompt the emergence, via mutation, of new phenotypes promoting long-term dispersion between patches and the evolution of policing mechanisms, both of which never would have had the opportunity to evolve in conditions where cells are free living (no scaffold regime). We could further imagine that this part of the environment becomes an integral part of the collective in the same way the extracellular matrix of multicellular organisms or the shell of molluscs—which has been literally internalized during evolution in some cephalopods (for details, see Kröger et al. [2011])—are considered integral parts of those organisms. In sum, a niche construction perspective is both compatible with the endogenization model presented in Section 5 and permits a more subtle explanation than one that regards the distinction between the internalist and the externalist modes of explanation I distinguished in Section 3 as referring to mutually exclusive categories.[12]

Before concluding, I should say a few words about the extent to which the Darwinian properties evolved under the reverted regime can be regarded as genuine Darwinian properties. I argued in Section 2 that for Lewontin's three conditions to correspond to Darwinian properties, they should be resilient to perturbation. A natural question arising from this remark is how much change is sufficient for the properties to be considered fully Darwinian rather than Darwinian-like. This type of question leads us into the territory of vagueness and sorites paradoxes (see Sorensen [2018]), which accept no definite answer. Even metazoans like us, which are regarded as paradigmatic examples of multicellular organisms, can potentially revert to more primitive forms of multicellularity or even unicellular organisms. Indeed, this is one hypothesis to explain the existence of cancer (Davies and Lineweaver [2011]; Vincent [2012]; Trigos et al. [2017]). For a broader context, see (Plutynski [2019]). It is well known that some types of cancer can be induced by changes in the environment, such as cancer induced by radiation (Shah et al. [2012]). However, there is also evidence that a cancerous cell can be 'normalized' (that is, made non-cancerous) when put in healthy tissue (Soto and Sonnenschein

---

[12] In addition to internalist and externalist explanations, Godfrey-Smith ([1996]) calls constructive explanations those explanations that explain external properties by invoking the activities of the objects under study.

[2011]). This points towards the view that considering genuine Darwinian properties as properties that would be resilient to all possible perturbations sets too high a bar.

Regarding the model presented in Section 5, that collectives do not revert to particles while they would have had there been no scaffolding initially is clearly a step towards making these evolved collectives Darwinian individuals in their own right. However, some classes of environmental changes would cause reversion to the unicellular state. Such would be the case, for instance, if a new high-growth mutant was produced in a single mutation (this can be observed in the NetLogo version of the model provided as supplementary material).[13] Such a new mutant cell would be part of the environment of other cells and would cause the demise of collectives. It is at this point that policing mechanisms—that is, mechanisms eliminating variation within collectives as proposed by Clarke—could render collectives more resilient to perturbation. Policing mechanisms prevent what Sober and Wilson ([1998]) have called 'subversion from within'.[14] The evolution of such mechanisms has not been the focus of this article, but studying their evolution is a natural extension of this work. Finally, note that resilience to changes in the environment could be obtained if collectives produce changes in the environment that render the environment inhospitable for single cells or give an advantage to patches that contain cells with a low growth rate, which leads us back to the potential role of niche construction in ETIs.

8 Conclusion

In this article, I proposed a model of ETIs: the evolution of Darwinian properties at a new level of organization. I showed that a population of particles can evolve more fully Darwinian properties from the existence of an ecological scaffold and a genotype with pleiotropic effects. An ecological scaffold permits the evolution of Darwinian-like properties at the collective level, while the existence of pleiotropic effects on dispersion over short and long timescales as well as growth rate, causes the endogenization of the scaffold and, consequently, the evolution of more fully Darwinian properties at the collective level. Thus, it fosters the evolution of collective-

---

[13] Recall that in the model, I have assumed that mutations are small. Consequently, there is no possibility that a high-growth mutant would be produced in a single or a few cell generations from a low-growth variant.

[14] Some attribute this term to Dawkins. While Dawkins's ([1976], p. 72) idea is the same, he uses the more agential term 'treachery from within'.

level properties that are insensitive to mutations that would otherwise lead to their demise. This result is significant in that it extends the ecological scaffolding model, making it a more likely scenario for the origins of ETIs, and provides new ways to explore the origins of multicellularity beyond those previously envisaged.

Acknowledgements

I would like to thank the members of the Theory and Method in Biosciences group both at the University of Sydney and Macquarie University for comments on a previous draft. In particular, I thank Guilhem Doulcier for his comments related to modelling. I am also thankful to Andrew Black, Ellen Clarke, Cody Nitschke, and Paul Rainey for discussions on the topic of this article. I gratefully acknowledge the financial support of the John Templeton Foundation (grant no. 62220). The opinions expressed in this article are those of the authors and not those of the John Templeton Foundation. I am the recipient of an Australian Research Council Discovery Early Career Research Award (Grant ID: DE210100303) funded by the Australian government.

Department of Philosophy
Macquarie University
North Ryde, NSW, Australia
and
Department of Philosophy & Charles Perkins Centre
University of Sydney
Sydney, NSW, Australia
p.bourrat@gmail.com

References

Abrams, M. [2012]: 'Measured, Modeled, and Causal Conceptions of Fitness', *Frontiers in Genetics*, **3**, p. 196.

Ågren, J. A. [2021]: *The Gene's-Eye View of Evolution*, Oxford: Oxford University Press.

Black, A. J., Bourrat, P. and Rainey, P. B. [2020]: 'Ecological Scaffolding and the Evolution of Individuality', *Nature Ecology and Evolution*, **4**, pp. 426–36.

Bourke, A. F. [2011]: *Principles of Social Evolution*, Oxford: Oxford University Press.

Bourrat, P. [2014]: 'From Survivors to Replicators: Evolution by Natural Selection Revisited', *Biology and Philosophy*, **29**, pp. 517–38.

Bourrat, P. [2015a]: 'Levels of Selection Are Artefacts of Different Fitness Temporal Measures', *Ratio*, **28**, pp. 40–50.

Bourrat, P. [2015b]: 'Levels, Time, and Fitness in Evolutionary Transitions in Individuality', *Philosophy and Theory in Biology*, **7**.

Bourrat, P. [2019]: 'Evolutionary Transitions in Heritability and Individuality', *Theory in Biosciences*, **138**, pp. 305–23.

Bourrat, P. [2021a]: *Facts, Conventions, and the Levels of Selection*, Cambridge: Cambridge University Press.

Bourrat, P. [2021b]: 'Transitions in Evolution: A Formal Analysis', *Synthese*, **198**, pp. 3699–731.

Bourrat, P., Doulcier, G., Rose, C. J., Rainey, P. B. and Hammerschmidt, K. [unpublished]: 'Beyond Fitness Decoupling: Tradeoff-Breaking during Evolutionary Transitions in Individuality', available at <www.biorxiv.org/content/10.1101/2021.09.01.458526v1.full>.

Bourrat, P. and Griffiths, P. E. [2018]: 'Multispecies Individuals', *History and Philosophy of the Life Sciences*, **40**, available at <doi.org/10.1007/s40656-018-0194-1>.

Calcott, B. and Sterelny, K. [2011]: *The Major Transitions in Evolution Revisited*, Cambridge, MA: MIT Press.

Caporael, L. R., Griesemer, J. R., Wimsatt, W. C., Müller, G. B., Schäfer, K. and Pradeu, T. [2013]: *Developing Scaffolds in Evolution, Culture, and Cognition*, Cambridge, MA: MIT Press.

Clarke, E. [2010]: 'The Problem of Biological Individuality', *Biological Theory*, **5**, pp. 312–25.

Clarke, E. [2013]: 'The Multiple Realizability of Biological Individuals', *The Journal of Philosophy*, **110**, pp. 413–35.

Clarke, E. [2014]: 'Origins of Evolutionary Transitions', *Journal of Biosciences*, **39**, pp. 303–17.

Clarke, E. [2016a]: 'A Levels-of-Selection Approach to Evolutionary Individuality', *Biology and Philosophy*, **31**, pp. 893–911.

Clarke, E. [2016b]: 'Levels of Selection in Biofilms: Multispecies Biofilms Are Not Evolutionary Individuals', *Biology and Philosophy*, **31**, pp. 191–212.

Damuth, J. and Heisler, I. L. [1988]: 'Alternative Formulations of Multilevel Selection', *Biology and Philosophy*, **3**, pp. 407–30.

Davies, P. and Lineweaver, C. H. [2011]: 'Cancer Tumors as Metazoa 1.0: Tapping Genes of Ancient Ancestors', *Physical Biology*, **8**, p. 015001.

Dawkins, R. [1976]: *The Selfish Gene*, Oxford: Oxford University Press.

De Monte, S. and Rainey, P. B. [2014]: 'Nascent Multicellular Life and the Emergence of Individuality', *Journal of Biosciences*, **39**, pp. 237–48.

DiFrisco, J. [2019]: 'Kinds of Biological Individuals: Sortals, Projectibility, and Selection', *British Journal for the Philosophy of Science*, **70**, pp. 845–75.

Doulcier, G., Hammerschmidt, K. and Bourrat, P. [unpublished]: 'Division of Labor and Fitness Decoupling in Evolutionary Transitions in Individuality', available at <www.biorxiv.org/content/10.1101/2020.10.23.352948v1.full>.

Doulcier, G., Lambert, A., De Monte, S. and Rainey, P. B. [2020]: 'Eco-evolutionary Dynamics of Nested Darwinian Populations and the Emergence of Community-Level Heredity', *eLife*, **9**, available at <doi.org/10.7554/eLife.53433>.

Garson, J. [2019]: *What Biological Functions Are and Why They Matter*, Cambridge: Cambridge University Press.

Glymour, B. [2017]: 'Cross-Unit Causation and the Identity of Groups', *Philosophy of Science*, **84**, pp. 717–36.

Godfrey-Smith, P. [1996]: *Complexity and the Function of Mind in Nature*, Cambridge: Cambridge University Press.

Godfrey-Smith, P. [2007]: 'Conditions for Evolution by Natural Selection', *Journal of Philosophy*, **104**, p. 489.

Godfrey-Smith, P. [2009]: 'Abstractions, Idealizations, and Evolutionary Biology', in A. Barberousse, M. Morange and T. Pradeu (*eds*), *Mapping the Future of Biology*, Dordrech: Springer, pp. 47–56.

Godfrey-Smith, P. [2011]: 'Darwinian Populations and Transitions in Individuality', in B. Calcott and K. Sterelny (*eds*), *The Major Transitions in Evolution Revisited*, Cambridge, MA: MIT Press, pp. 65–81.

Griesemer, J. R. [2000]: 'The Units of Evolutionary Transition', *Selection*, **1**, pp. 67–80.

Griesemer, J. R. [2005]: 'The Informational Gene and the Substantial Body: On the Generalization of Evolutionary Theory by Abstraction', in M. R. Jones and N. Cartwright (*eds*), *Idealization XII: Correcting the Model—Idealization and Abstraction in the Sciences*, Amsterdam: Rodopi, pp. 59–115.

Grosberg, R. K. and Strathmann, R. R. [2007]: 'The Evolution of Multicellularity: A Minor Major Transition?', *Annual Review of Ecology, Evolution, and Systematics*, **38**, pp. 621–54.

Guay, A. and Pradeu, T. [2015]: *Individuals across the Sciences*, Oxford: Oxford University Press.

Haig, D. [2012]: 'The Strategic Gene', *Biology and Philosophy*, **27**, pp. 461–79.

Hammerschmidt, K., Rose, C. J., Kerr, B. and Rainey, P. B. [2014]: 'Life Cycles, Fitness Decoupling, and the Evolution of Multicellularity', *Nature*, **515**, pp. 75–79.

Jablonka, E. and Lamb, M. J. [2006]: 'The Evolution of Information in the Major Transitions', *Journal of Theoretical Biology*, **239**, pp. 236–46.

Kröger, B., Vinther, J. and Fuchs, D. [2011]: 'Cephalopod Origin and Evolution: A Congruent Picture Emerging from Fossils, Development, and Molecules: Extant Cephalopods Are Younger than Previously Realised and Were under Major Selection to Become Agile, Shell-Less Predators', *BioEssays*, **33**, pp. 602–13.

Laland, K., Matthews, B. and Feldman, M. W. [2016]: 'An Introduction to Niche Construction Theory', *Evolutionary Ecology*, **30**, pp. 191–202.

Lehmann, L. [2007]: 'The Evolution of Trans-generational Altruism: Kin Selection Meets Niche Construction', *Journal of Evolutionary Biology*, **20**, pp. 181–89.

Lehmann, L. [2008]: 'The Adaptive Dynamics of Niche Constructing Traits in Spatially Subdivided Populations: Evolving Posthumous Extended Phenotypes', *Evolution*, **62**, pp. 549–66.

Lewontin, R. C. [1970]: 'The Units of Selection', *Annual Review of Ecology and Systematics*, **1**, pp. 1–18.

Lewontin, R. C. [1983]: 'The Organism as the Subject and Object of Evolution', *Scientia*, **77**, p. 65.

Lewontin, R. C. [1985]: 'Adaptation', in R. Levins and R. C. Lewontin (*eds*), *Dialectics and Reductionism in Ecology*, Cambridge, MA: Harvard University Press, pp. 65–84.

Libby, E., Conlin, P. L., Kerr, B. and Ratcliff, W. C. [2016]: 'Stabilizing Multicellularity through Ratcheting', *Philosophical Transactions of the Royal Society B*, **371**, available at <doi.org/10.1098/rstb.2015.0444>.

Libby, E. and Ratcliff, W. C. [2014]: 'Ratcheting the Evolution of Multicellularity', *Science*, **346**, pp. 426–27.

Lidgard, S. and Nyhart, L. K. [2017]: *Biological Individuality: Integrating Scientific, Philosophical, and Historical Perspectives*, Chicago, IL: University of Chicago Press.

Lu, Q. and Bourrat, P. [2018]: 'The Evolutionary Gene and the Extended Evolutionary Synthesis', *British Journal for the Philosophy of Science*, **69**, pp. 775–800.

Maynard Smith, J. and Szathmary, E. [1995]: *The Major Transitions in Evolution*, Oxford: Oxford Univeristy Press.

Michod, R. E. [1999a]: *Darwinian Dynamics*, Princeton, NJ: Princeton University Press.

Michod, R. E. [1999b]: 'Individuality, Immortality, and Sex', in L. Keller (*ed.*), *Levels of Selection in Evolution*, Princeton, NJ: Princeton University Press, pp. 53–74.

Michod, R. E. [2005]: 'On the Transfer of Fitness from the Cell to the Multicellular Organism', *Biology and Philosophy*, **20**, pp. 967–87.

Michod, R. E. and Roze, D. [1997]: 'Transitions in Individuality', *Proceedings of the Royal Society of London B*, **264**, pp. 853–57.

Michod, R. E., Viossat, Y., Solari, C. A., Hurand, M. and Nedelcu, A. M. [2006]: 'Life-History Evolution and the Origin of Multicellularity', *Journal of Theoretical Biology*, **239**, pp. 257–72.

Odling-Smee, F. J., Laland, K. N. and Feldman, M. W. [2003]: *Niche Construction: The Neglected Process in Evolution, Monographs in Population Biology*, Princeton, NJ: Princeton University Press.

Okasha, S. [2005]: 'Multilevel Selection and the Major Transitions in Evolution', *Philosophy of Science*, **72**, pp. 1013–25.

Okasha, S. [2006]: *Evolution and the Levels of Selection*, Oxford: Oxford University Press.

Okasha, S. [2016]: 'The Relation between Kin and Multilevel Selection: An Approach Using Causal Graphs', *British Journal for the Philosophy of Science*, **67**, pp. 435–70.

Okasha, S. [2021]: 'The Strategy of Endogenization in Evolutionary Biology', *Synthese*, **198**, pp. 3413–35.

Pepper, J. W. and Herron, M. D. [2008]: 'Does Biology Need an Organism Concept?', *Biological Reviews*, **83**, pp. 621–27.

Plutynski, A. [2019]: 'Cancer', in E. N. Zalta (*ed.*), *The Stanford Encyclopedia of Philosophy*, available at <plato.stanford.edu/archives/sum2019/entries/cancer>.

Pradeu, T. [2012]: *The Limits of the Self: Immunology and Biological Identity*, Oxford: Oxford University Press.

Rainey, P. B. and Kerr, B. [2010]: 'Cheats as First Propagules: A New Hypothesis for the Evolution of Individuality during the Transition from Single Cells to Multicellularity', *BioEssays*, **32**, pp. 872–80.

Rainey, P. B. and Rainey, K. [2003]: 'Evolution of Cooperation and Conflict in Experimental Bacterial Populations', *Nature*, **425**, pp. 72–74.

Rose, C. J., Hammerschmidt, K. and Rainey, P. B. [2020]: 'Meta-population Structure and the Evolutionary Transition to Multicellularity', *Ecology Letters*, available at <doi.org/10.1111/ele.13570>.

Scott-Phillips, T. C., Laland, K. N., Shuker, D. M., Dickins, T. E. and West, S. A. [2014]: 'The Niche Construction Perspective: A Critical Appraisal', *Evolution; International Journal of Organic Evolution*, **68**, pp. 1231–43.

Shah, D. J., Sachs, R. K. and Wilson, D. J. [2012]: 'Radiation-Induced Cancer: A Modern View', *The British Journal of Radiology*, **85**, pp. e1166–73.

Shelton, D. E. and Michod, R. E. [2014]: 'Group Selection and Group Adaptation during a Major Evolutionary Transition: Insights from the Evolution of Multicellularity in the Volvocine Algae', *Biological Theory*, **9**, pp. 452–69.

Shelton, D. E. and Michod, R. E. [2020]: 'Group and Individual Selection during Evolutionary Transitions in Individuality: Meanings and Partitions', *Philosophical Transactions of the Royal Society B*, **375**, available at <doi.org/10.1098/rstb.2019.0364>.

Sober, E. [1984]: *The Nature of Selection*, Cambridge, MA: MIT Press.

Sober, E. and Wilson, D. S. [1998]: *Unto Others: The Evolution and Psychology of Unselfish Behavior*, Cambridge, MA: Harvard University Press.

Sorensen, R. [2018]: 'Vagueness', in E. N. Zalta (*ed.*), *The Stanford Encyclopedia of Philosophy*, available at <plato.stanford.edu/archives/sum2018/entries/vagueness/>.

Soto, A. M. and Sonnenschein, C. [2011]: 'The Tissue Organization Field Theory of Cancer: A Testable Replacement for the Somatic Mutation Theory', *BioEssays*, **33**, pp. 332–40.

Sterner, B. [2015]: 'Pathways to Pluralism about Biological Individuality', *Biology and Philosophy*, **30**, pp. 609–28.

Tanaka, M. M., Godfrey-Smith, P. and Kerr, B. [2020]: 'The Dual Landscape Model of Adaptation and Niche Construction', *Philosophy of Science*, **87**, pp. 478–98.

Trigos, A. S., Pearson, R. B., Papenfuss, A. T. and Goode, D. L. [2017]: 'Altered Interactions between Unicellular and Multicellular Genes Drive Hallmarks of Transformation in a Diverse Range of Solid Tumors', *Proceedings of the National Academy of Sciences USA*, **114**, pp. 6406–411.

van Gestel, J. and Tarnita, C. E. [2017]: 'On the Origin of Biological Construction, with a Focus on Multicellularity', *Proceedings of the National Academy of Sciences USA*, **114**, pp. 11018–26.

Veit, W. [2021]: 'Scaffolding Natural Selection', *Biological Theory*, available at <doi.org/10.1007/s13752-021-00387-6>.

Vincent, M. [2012]: 'Cancer: A De-repression of a Default Survival Program Common to All Cells?', *BioEssays*, **34**, pp. 72–82.

Wilensky, U. [1999]: 'NetLogo', Center for Connected Learning and Computer-Based Modeling, Northwestern University, available at <ccl.northwestern.edu/netlogo/>.

Williams, G. C. [1966]: *Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought*, Princeton, NJ: Princeton University Press.

Wright, S. [1931]: 'Evolution in Mendelian Populations', *Genetics*, **16**, pp. 97–159.