
Origins of evolutionary transitions

ELLEN CLARKE

All Souls College, 27 High Street, Oxford OX1 4AL, UK

(Email, ellen.clarke@all-souls.ox.ac.uk)

An ‘evolutionary transition in individuality’ or ‘major transition’ is a transformation in the hierarchical level at which natural selection operates on a population. In this article I give an abstract (i.e. level-neutral and substrate-neutral) articulation of the transition process in order to precisely understand how such processes can happen, especially how they can get started.

[Clarke E 2014 Origins of evolutionary transitions. *J. Biosci.* **39** 303–317] DOI 10.1007/s12038-013-9375-y

1. Introduction

Most existing work on major transitions in evolution takes as its primary explanandum the success of cooperative aggregation in the face of free-rider problems. There are, however, several other necessary components of an explanation of evolutionary transitions. I define an evolutionary transition as a shift in the hierarchical level at which heritable fitness variance is expressed in a population, before distinguishing different kinds of questions that can be posed with respect to transitions and then setting out the conceptual ingredients required to answer each kind. Then I describe a ‘chicken and egg problem’, which threatens attempts to explain evolutionary transitions in individuality, and demonstrate how the problem can be resolved. Finally I describe the details that a successful selective account of an evolutionary transition ought to include. The article is a work of conceptual analysis: that is, the conclusions are drawn, not from examination of empirical data, but only via reflection on existing biological theory.¹

2. What is a transition?

‘Major transition in evolution’ is a collective name for a group of events which have happened at different points in evolutionary history, during which natural selection transformed formerly free-living organisms into mere parts in new, higher-level wholes. The modern debate was framed by Maynard Smith and Szathmáry (1995) although the basic ideas

were developed in earlier works by John Tyler Bonner (1974) and Leo Buss (1987), and have subsequently been developed by Elliott Sober and David Sloan Wilson (1998), Rick Michod (1999), Samir Okasha (2005, 2006) and Peter Godfrey Smith (2009). One of Maynard Smith and Szathmáry’s contributions was to define a major transition by identifying a pattern that is common across an otherwise diverse set of events:

‘Entities that were capable of independent replication before the transition can replicate only as part of a larger whole after it’ (Maynard Smith and Szathmáry 1995 p 8).

For example, mitochondria, the energy-producing organelles found in all eukaryote cells, can replicate only as parts of eukaryotic cells, but they evolved from ancestors which were capable of independent replication. The transformation from a unicellular into a multicellular state is one of the paradigmatic, and best-studied, kinds of evolutionary transition that are thought to have taken place. It is estimated that the transition to multicellularity has taken place at least 25 times – in other words, 25 different cellular lineages have independently made the jump to communally organized life (Bourke 2011). Other examples include: the origin of the cell, in which replicating molecules became aggregated into compartments containing many molecules, about 3500 million years ago; the emergence of chromosomes from independent replicators; the creation of eukaryotes by endosymbiosis; and the transition of multicellulars to colonial life, for example in eusocial insects, around 150 million years ago. None of these transformations was inevitable – all of the

Keywords. Evolutionary transition; individuality; level of selection

¹ This is a philosophical strategy, intended to complement, rather than to replace, the more standard strategy of generalizing from data.

original individuals have non-aggregated descendants, as well as aggregated descendants.

Maynard Smith and Szathmáry included several additional events, which I have omitted here because they are more controversial. They included the origin of sex, and also of language, because they characterized the transitions as being a reorganization of the way in which information is transmitted across generations. However, most authors now prefer a more structurally compositional definition. On a compositional view, what unifies the different transitions is the process of wholes becoming parts that are physically nested inside new higher-level wholes. A major transition is a process in which separate parts become lumped together.

In this article I follow Buss in adopting a selective definition – according to which a major transition consists in a transformation of the hierarchical level at which selection operates on a population.² This implies that a mere aggregation of entities into groups is insufficient (Mayr 1982; Newman 2003). The entities need to be compounded in such a way that higher-level selection takes place. There have been many major evolutionary events that this definition of a transition excludes. The evolution of flight, for example, constituted a step-change in bird evolution, but not one that involved any composition of birds into higher-level entities. The origin of segmentation, for another, involves a new relationship between parts and whole, but does not constitute a major transition, on a selective definition, because there is no change in the level at which selective competition occurs.

My aim is to make the abstract logic of the transition process as precise as possible, in order to obtain clarity on questions about how they happen. There are two qualitatively distinguishable states, which we can call State One and State Two. A State One population is divided into objects with the following properties – they are living, and they exhibit heritable variance in fitness (figure 1). To understand the second claim, assign each object in the population a character value or trait, z , and a fitness, w . The population must meet the following three conditions:

- i. Character value (z) must vary so that different members of the population express different values for z (phenotypic variation).
- ii. Variation in z must correlate with variation in fitness, w (differential fitness).
- iii. Fitness (w) must be heritable (fitness is heritable).

² Note that I am not committed to saying that selection *causes* transitions – in fact I will deny that selection can bring about the first transitional step. Instead I am using a definition on which a change in the level of selection is what *constitutes* a transition. I defend this view in Clarke 2013, arguing that only a selectively defined concept is capable of underpinning the successful inferences we achieve regarding the action of natural selection.




Figure 1. A State One population.

Together these conditions guarantee that the State One population is capable of undergoing evolution by natural selection (Lewontin 1970). The objects in a State One population therefore compete against each other in a standard one-level selection process. We call the objects in State One population ‘organisms’ or sometimes ‘biological individuals’.

A State Two population is identical except that the objects into which it is divided, and which exhibit heritable variance in fitness, are themselves aggregates of former State One objects (figure 2).³ These aggregate creatures can also be called ‘organisms’, but their parts cannot. For precision, I add the further constraint that State Two populations exhibit zero heritable variance in fitness at the level of the parts (as in Gardner and Grafen’s 2009 ‘superorganism’).⁴ In other words, heritable fitness variance is exhibited exclusively at one hierarchical level, in each case.

This distinction between different states that a population can be in is conceptually precise, as well as intuitively familiar.⁵ We tend to think of unicellular life forms as composing State One populations, and multicellular life forms as composing State Two populations. Then, a solitary

³ Note that the parts of State One and State Two objects may or may not be identical. In other words, a transition can be fraternal or egalitarian (Queller 2000). In real-life cases, the parts are unlikely to ever be genetically identical, and will sometimes be from different species. So there is nothing in this schema which precludes symbiotic or chimeric higher-level organisms.

⁴ This is an idealization. Real-life organisms will almost never meet this ideal – that is, there will always be *some* lower-level conflict, excepting perhaps very tiny organisms.

⁵ Note that sometimes we use the term ‘higher-level’ to refer to lower-level properties viewed from a particular scale. For example, water molecules do not exhibit waves but water does. This is different from my use, in which to say that selection is at the higher level is to make a claim about the scale at which variance occurs. This is not a matter of perspective – meiotic drive, for example, either takes place or it does not (Okasha 2006).




Figure 2. A State Two population.

amoeba might be an approximation of a State One object, while a polar bear is an approximation of a State Two object. However, real-life organisms need not fall into one category or other – many will be stable at an intermediate step.

We can define an evolutionary transition as consisting in the transformation of a State One population into a State Two population (figure 3).

The challenge of explaining a transition is not to explain how amoebae can be transformed into polar bears, or even how human muscle cells can be transformed into human beings, but rather, to understand how an object of the *type* exemplified in State One could be transformed into an object of the *type* exemplified in State Two. Given that transitions between State One organisms and State Two organisms have taken place over huge timescales,⁶ we do not have access to extant lineages containing cases from both states. The closest we can get is to compare current State Two objects with what we know of their lower-level ancestors from fossils, or to compare State Two objects with close relatives which we take to be similar to the lower-level ancestors of the objects. In this way, *Volvox carteri* are often conceived as exemplars of State Two objects in contrast with *Chlamydomonas reinhardtii*, their unicellular cousins (Kirk and Harper 1986; Larson *et al.* 1992; Koufopanou 1994; Kirk 1998; Michod 1999). We do not think that there is any sense in which *V. carteri* have transitioned from *C. reinhardtii*, however, any more than we think humans are descended from chimpanzees. But we do think that *V. carteri* have transitioned from an ancestor which, in respect of its unicellularity at least, was similar to *C. reinhardtii*.

⁶ This is not to imply that partial steps in a transition cannot take place relatively suddenly, but a whole transition – all the way from no selection at level *x* to exclusive selection at *x* – will probably require a multiplicity of steps, and will happen gradually, rather than all at once. This is not because the changes have to be genetic, but only because the elimination of lower-level selection is not a simple thing to accomplish, and usually requires several different complex policing mechanisms working simultaneously. See section 4.




Figure 3. A transition process.

Major transitions have been enormously significant events in history, which have repeatedly rewritten the evolutionary rulebooks by radically reorganizing the terms under which natural selection is played out. It is not therefore surprising that biologists and philosophers have been enraptured for some time now by the set of puzzles that understanding the major transitions presents. If we can explain the process of an evolutionary transition, then we can begin to understand one of the most powerful evolutionary phenomena the earth has ever witnessed. But what does explaining a transition involve? What key puzzles are there to be solved?

3. Different explanatory challenges

There are a number of different ongoing research questions associated with the major transitions (Bourke 2011; Calcott 2011), which we can organize into three classes.⁷ The first asks *how* transitions occur: what are the mutations or other changes that first get the process going and then drive it along? Ideally we would like to be able to list a series of mutational or developmental steps by which the properties that we associate with organismality can be reorganized so that they are expressed at a new, compositionally hierarchical level (Mayr 1982).

Secondly, we can ask questions about *why* transitions occur – for what reason does a population sometimes change from State One to State Two? In other words, what sort of fitness benefits are made available by moving to a higher level of organization? Calcott calls this the problem of ‘generating benefit’ (Calcott 2008). What is it about the proximate mutations/novel traits that makes them selectively advantageous? To answer these questions we need to be realistic, and consider all the connotations of a shift to life lived in close proximity to former competitors. The benefits wrought will have to outweigh the costs incurred.

⁷ Bourke divides the process of transition itself into three distinct phases – social group origination, maintenance and transformation (Bourke 2011 p 15). I prefer to think of the transition process as smoothly continuous rather than carving it up into discrete segments, but nonetheless agree that there are distinguishable questions to be asked about it.

Finally, we can ask questions about how the higher-level organization is *maintained*, especially what protects cooperative interactions among the parts of a new higher-level organism from invasion by free-riding cheats. What makes the spread of the novel traits robust? In virtue of what are the objects in State Two evolutionarily stable?

An overwhelming focus in the transitions literature thus far has been on the third question, which has been framed thus

‘Why did not natural selection, acting on entities at the lower level, disrupt integration at the higher level?’
(Maynard Smith and Szathmáry 1995 p 8)

Communal life usually comes at a cost – the loss of a freedom, the production of a common good, or just the tight sharing of space. The puzzle is, why do not members tend to evolve selfish mutations that allow them to reap the rewards of higher-level organization, but without shouldering the costs? Another way to put the problem is that after the new collectivist life forms have spread through the population, the situation – the terms of the evolutionary game – have changed. Where initially there were two possible strategies – live together or die alone – a third player, ‘cheat’, can now enter the game. The aggregate forms might now have to compete against the cheats as well as against the solitary ancestors. Implicit is the thought that there will always be a way to cheat – to take advantage of the aggregates – and that some mutation or other will always find it.

Buss characterized this problem as one of explaining how selection for a higher-level property can come to overwhelm lower-level⁸ selective pressures (Buss 1987).⁹ Most of the debate now concerns the best conceptualization of the solution to this problem (Bull and Rice 1991; Kerr and Godfrey-Smith 2002; Sachs *et al.* 2004; Lehmann *et al.* 2007; Fletcher and Doebeli 2010; Gardner *et al.* 2011; Marshall 2011; Van Veelen *et al.* 2012). At a very general level, there is broad consensus that the solution to the free-rider problem consists in there being some sort of guarantor of assortment among cooperating particles, so that they are more likely to interact with one another than with cheats (Hamilton 1971; Fletcher and Doebeli 2009; Bijma and Aanen 2009). Much progress has been made in identifying various kinds of ‘policing’ mechanisms that have evolved for the purpose of raising assortment by mediating or eliminating conflicts between the parts of a collective. For example, honeybee workers monitor and destroy illicit egg-laying by their fellows (Ratnieks 1988; Ratnieks *et al.* 2006). Developmental

bottlenecks sieve out genetic mutants from multicellular organisms to reduce the likelihood of conflicts breaking out among rival genotypes (Dawkins 1982). High relatedness among the parts of an organism aligns their evolutionary interests (Hamilton 1964). Immune surveillance ameliorates the problem of cheaters in symbiotic complexes (Gilbert *et al.* 2012 p 333). Most of the current debate now concerns the extent to which different mechanisms are sufficient in offering the relevant guarantee (Nowak 2006; West, Griffin and Gardner 2007; Foster 2009; Leimar and Hammerstein 2010).

Note, however, that the third maintenance question can only be asked in a context which assumes that questions one and two have answers (Calcott 2011; Trestman 2013). We have to take it for granted that there is a good explanation for how a transition first originated (i.e. that there are some plausible first steps that would cause aggregation of the right kind among the State One objects) *and* that those aggregates which took the first steps would proliferate, because it gives them some sort of advantage over their solitary competitors. The third question assumes all this and further that, even so, transitions should always fail. This tendency to blackbox questions about the origins of transitions is problematic, because these questions are interesting in their own right, and not as easy to answer as is perhaps assumed.¹⁰ Furthermore, the details of how the origin questions can be and have been solved, during each of the numerous transitions that have taken place in the history of life, surely have implications for the maintenance problem too. The way in which the first how and why questions got answered determines the nature of the aggregates in which policing must evolve. What is more, perhaps sometimes the relevant series of developmental steps just happens to be such that there is no plausible way to cheat. Perhaps sometimes the relevant evolutionary advantage is so great that the transition can be maintained even in the face of significant cheating. In this way, exploration of proximate and evolutionary questions about the early origins of evolutionary transitions has the power to transform, and also to undercut, the debates about how transitions, once initiated, are maintained.

4. The How-question: A first step

A major transition consists in the transformation of a State One population, in which selection acts exclusively at level *x*, into a State Two population, in which selection acts exclusively¹¹ at

⁸ It is important to be clear that when I say ‘object-level selection’ or ‘lower-level selection’, I always refer to variation in fitness among objects *within groups*, not to variation in fitness among objects in the global population (Sober 2011).

⁹ Buss’ own answer to the problem focuses on the role of germ soma separation in eliminating competition between different cell lineages in a multicellular.

¹⁰ For example, it may be assumed that reproducing clumps of cells are easy to generate, and offer obvious selective advantages, without giving proper consideration to issues such as how reproduction of the whole emerges, or how such clumps can overcome some obvious difficulties, such as the build-up of waste products.

¹¹ Again, this is idealized – in real-life transitions the start and end points are unlikely to be perfectly exclusive levels of selection.

level $x+1$. The transitions literature seeks answers to questions about how this happens, why it happens, and what prevents the process from being derailed by lower-level competition. Here I focus on the first question – how do State One populations become transformed into State Two populations? While others have focused on the end stages of such a process (Gardner and Grafen 2009), or have tried to identify a critical threshold or halfway point marking a difference between the two states (Buss 1987; Damuth and Heisler 1988; Sober and Wilson 1998; Michod 1999; Okasha 2006; Bourke 2011), here I pay particular attention to the very early stages of the process, asking how the trajectory is able to first get started.

The starting point of my analysis is consideration of a system in State One – a population of well-mixed Darwinian particles. The How-question demands a proximate, mechanistic account of the causes involved in moving such a population out of State One, and into a state that is intermediate between States One and Two. There are two parts to this task. One is to say, in a definitional sense, what kind of change constitutes a first step in a transition process. In other words, in virtue of what does a population in State One leave that state, and begin its journey towards State Two? The second is to give a causal how-possibly account of how such a first step might be taken – to identify mechanisms or properties which could plausibly appear and qualify an incipiently transitioning population, according to our definition. Meeting this second challenge involves getting concrete about the actual processes and entities involved in the early stages of a transition.

My focus will be initially on the first challenge: in exactly what circumstances do we want to say that a population has departed from State One? I seek a very general answer to this question, one that can encompass not only different proximate ways of taking the first step of any given transition but also first steps for any different kind of transition. For example, the objects in the population under consideration might be cells which compete with one another in a standard one-level selection process. But equally, they might be solitary multicellulars, or even pre-cellular replicating molecules of some kind. We need not expect the actual mechanisms or developmental steps to be identical in each case, but neither would this necessarily stand in the way of giving a common abstract characterization of what is significant about the step in each case. In this way the current analysis departs from those offered by Kirk (2005), Michod (1999) and others (Rainey and Rainey 2003; Grosberg and Strathmann 2007; Rokas 2008; Sachs 2008), because theirs were intended specifically to describe the first steps of a transition to multicellularity (in a specific clade, in Kirk and Michod's cases) while here I seek a fully level-neutral account of transition.

So, what does it mean to leave State One? Returning to the definitions I gave earlier, to be an object in this state

requires the following qualifications: being alive; exhibiting heritable variance in fitness exclusively at the between-object level. Death is an obvious route to exiting the state, but since the life/death distinction is not the one I am trying to capture here, let us assume it is not taken. Then, the population leaves State One just in case it fails to exhibit heritable fitness variance entirely, or in case it fails to exhibit such variance *exclusively* at the between-object level. The latter condition will hold just in case some degree of heritable fitness variance gets expressed at either a lower level – among parts of the State One objects – or at a higher level – between aggregates of State One objects. In the former case, we could say that the State One population has begun to transition downwards, towards a lower level of organization. The latter case fits the more paradigmatic vision of an evolutionary transition, as journeying upwards, towards a higher level of organization. I will focus on this type of case, but the other is analogous.

According to this view, a population takes a first step out of State One as soon as some degree of heritable variance in fitness shows up at the level of aggregates, as well as, or instead of, at the level of objects. I assume that transitions occur gradually, rather than in discrete jumps (see why in section 4), so that the former is likely to be temporally prior. In other words, we should expect that the first initial move is to a situation where only a small degree of fitness variance is manifested at the higher level, simultaneously with a large amount of variance being manifested at the object level. In State One, selection acts exclusively at the lower level, because this is the only level at which there is the right sort of fitness variance. In State Two, selection acts exclusively at the higher level, again, because this is the only level at which objects vary in the right way. Between these two extremes we should expect a continuum of intermediate possibilities, of populations in which some fitness variance appears at one level, *simultaneously with variance at the other level*. As we follow the gradient of this continuum, we will see that the *degree* to which the objects in the population exhibit heritable variance in fitness as aggregates, rather than individually, incrementally increases.

Note that while Buss and others define transitions in terms of actual selection, I talk about heritable fitness variance instead. Since that variance acts as the raw material on which selection acts, my definition avoids problems stemming from the fact that actual selection-at-a-level may fail to occur altogether, for reasons that are distinct from whether or not a population has transitioned or not – for example, population size or environmental contingencies. The existence of heritable fitness variance gives a population the capacity to undergo selection, whether or not that capacity is actually realized. So my definition enables us to get a handle on the essential question – what needs to become true of a population of State One objects before higher-level selection is able

to operate on the population? The important point is that by shifting heritable fitness variance up to a higher hierarchical level, a first transitional step makes *possible* the action of higher-level selection.

Two questions make themselves apparent at this stage – First, how can we know when some fitness variance has appeared at the higher level? Second, what could bring about such a change?

One way to answer the first question is by applying a statistical measure of higher-level selection. Price's multi-level covariance equation (equation 1) provides a popular way of measuring the strength of selection at different hierarchical levels (Price 1970, 1972) and works well in this context, although there are many alternative measures that could be used in its place.¹²

$$\bar{w}\Delta\bar{z} = \left(E_k \left(\text{Cov}_k(w_{jk}, z_{jk}) \right) + \text{Cov}(W_k, Z_k) \right) + E(w_i \Delta z_i)$$

**average object-level
covariance**
**group-level
covariance**
**transmission
bias**

↖
↗
↑

Equation 1. Here z denotes the trait value of an object, \bar{z}_i is the average trait value in the whole population, and z_{jk} denotes the trait value of the j^{th} object in the k^{th} group. So 'k' indexes groups, while 'j' indexes objects within groups, and 'i' indexes objects in the global population. Fitness w is defined in the corresponding manner. Z_k denotes the trait value of the k^{th} group, defined as the average trait value of its constituent objects, and W_k denotes the fitness of the k^{th} group, defined as the average fitness of its constituent objects. Cov denotes covariance and E denotes expectation, or average (taken from Okasha 2006 pp 62–64). Price showed that the population-level change in the value of a trait over one generation can be decomposed into three terms (Price 1972). The 'expectation term' – the first term on the right hand side – reflects the average of the within-group covariances between object trait value (z) and object fitness (w). The 'covariance term' – the second term on the right hand side – reflects covariance between the trait value expressed by *groups* (Z) and *their* fitness (W). Finally, the 'transmission bias', the final term, reflects change that is due to imperfect heritability in object replication (see also Frank 2012b).

The single-level Price equation describes the change in the population average value of a trait over one generation as a sum of the change that is due to selection, and the change that is due to transmission bias. The change due to selection is equal to the statistical covariance between the fitnesses w and the trait values z of the objects in the population.¹³ The multilevel Price equation, applied to a population of objects organized into groups, breaks the global covariance term down into two further components that are more and less local to particular groups. The first part is the average of the within-group covariances between object trait value and object fitness. To calculate this we treat each group in turn, calculating the local covariance between trait value and fitness of objects, before averaging the results for all the groups. The second part is the covariance between the trait expressed by *groups* and *their* fitness. Note that the equation itself does *not* tell us

anything about how to choose these groups – they must be defined before the multilevel analysis can be applied – I will say more about this shortly. The fitness of a group can be understood as the average of the member objects' fitnesses (this has been called 'MLS1 fitness'; Damuth and Heisler 1988; Okasha 2001) or alternatively as the number of daughter groups produced ('MLS2 fitness'). It makes no difference to the Price formalization which of these we use (Okasha 2006). People find it natural to think of the 'expectation term' (the first term on the right hand side) as describing the change that is due to selection within groups, while the 'covariance term' (the second term on the right hand side) describes the change that is due to selection *between* groups.

On a Pricean view, we would treat a transition as having been initiated as soon as the between-group component (the second term on the right hand side) of the multilevel Price equation takes a non-zero value. Given the analytic equivalence of Lewontin's conditions and Price's equation (Okasha 2006 p 36), we should understand this as a sort of restatement of the earlier claim, that a transition begins as soon as heritable variance in fitness is first manifested at the level of aggregates, rather than by lower-level objects alone. Price's equation is not intended to reveal anything about the

¹² For example, contextual analysis (Heisler and Damuth 1987; Goodnight *et al.* 1992; Goodnight 2013); neighbourhood analysis (Nunney 1985; Nowak and May 1992; Godfrey-Smith 2008); inclusive fitness analysis (Hamilton 1964; Gardner, West and Wild 2011).

¹³ Taken from Okasha 2006. See also Gardner 2008.

underlying causes of selection – about what properties the objects need to have in order for a covariance between character values and fitness to emerge. Nothing is filled in about what mechanisms are sufficient to correlate character value, z with fitness, w , or about what will guarantee variation in z , and so on. The Price formalization simply helps us to see that these are the conditions that must be met, by whatever means. Both perspectives (Price's and Lewontin's) are outcome-based rather than offering any mechanistic details about how the change might come about. For the latter detail, we need to move to task two – to a description of actual mechanisms which can bring about the relevant outcome.

However, there are some further conceptual matters that have to be settled first. The Price approach points to the exhibition of higher-level heritable fitness variance whenever the second term of the multilevel equation is non-zero. Much as a thermometer will register a change in room temperature without implying anything about the causes of that change, Price's equation can detect the manifestation of higher-level fitness variance, but is not in any way informative about *why* variance should be appearing at that level (Frank 2012a). The equation does not distinguish among processes underlying a given change in trait frequencies. Some scepticism about the Price approach (Nowak and Highfield 2011; Van Veelen *et al.* 2012) is motivated by the fact that statistical correlation may or may not be the product of causation. In order to eliminate those cases in which higher-level fitness variance shows up as a mere epiphenomenon or statistical artefact (Williams 1966) of lower-level variance (a 'cross-level by-product' Okasha 2006), we need to think about the second question, and identify possible causes of higher-level (within-group) fitness correlation. This issue dovetails with that of how to define the groups in the multilevel model.

We can generate within-group correlation simply by choosing to group State One objects on the basis of their fitness, assigning like-fitness with like-fitness. This method could generate a population structure in which all of the fitness variance occurs exclusively at the between-group level, just because we have defined the groups so that every member's fitness is identical to that of its group-mates. This is an example of how higher-level variance can show up as a mere statistical artefact of lower-level variance, because the result is explained entirely by the modeller's choice of groupings. We need to rule such cases out in order to use the covariance approach as a method for uncovering true underlying causality. This problem is not as complex as it might appear, however. All we need to do is restrict ourselves to consideration of populations in which fitness correlations are being caused by some sort of biological property or mechanism, rather than being imposed on the population by ourselves. We do not need to get too worried about the definition of biological mechanism – the essential point is merely that something in the world itself, rather than

in the modeller's choice of representation, is sorting objects into groups, in such a way that the fitness of one object is more predictive of the fitness of each of its group mates, on average, than it is of the fitness of a randomly chosen object from the rest of the population.

I will suggest some concrete examples of mechanisms that will do the necessary work of generating non-subjective groups in section 5, but first it is worth providing an abstract description of what is essential to such mechanisms. In other words, what does it mean to generate the right sort of group?

It is important that we do not merely resort to a common-sense definition of 'group', as physically stuck together in some way, or to a common-sense definition of 'organism', as being functionally differentiated, or genetically homogeneous and so on. The reason is that while these particular properties may be sufficient to bring about fitness correlation, in the right circumstances, it is the effect – the fitness assortment – that is essential to the first transitional step, *not* the particular manner in which it is brought about.¹⁴ Physical cohesion, for example, seems like a fairly universal element of the transition to multicellularity from single celled ancestors, which is why it is reasonable that researchers focusing on this transition tend to emphasize the significance of incomplete separation after mitosis as a critical first transitional step (Bonner 1998; Kirk 2003; Sachs 2008). But of course, if we think instead about the transition to eusociality in social insects, it is not helpful to posit a mechanism that could stick offspring insects to their mother. Instead, the focus is on continuity of behavioural interactions among the insects. In both cases, the highlighted change is relevant only in so far as it is causally associated with a correlation between the heritable fitnesses of the objects involved. We must avoid elevating proximate mechanisms into definitional criteria if we want to arrive at a truly level-neutral and general analysis of an evolutionary transition.

My answer is that groups are generated by mechanisms which cause fitness alignment amongst their member objects. This is an idea that has been expressed in many ways – as fate-coupling (Buss 1987), policing (Michod 1999; Okasha 2006), cooperation (Queller and Strassmann 2009) assortment (Fletcher and Doebeli 2009) correlation (Folse and Roughgarden 2010) and 'cohesion generation' (Hamilton *et al.* 2009). In other work I define the effect of any 'policing mechanism' as inhibiting the capacity of an aggregate to undergo within-aggregate selection (Clarke 2013). In section 5 I will say more about the variety of concrete processes and properties which will achieve this function. Fitness alignment is the essential outcome of any grouping mechanism,

¹⁴ After all, we could use mechanical means to stick together a group of humans, but would not therein consider the group to constitute a super-organism – why not?

because by tying together the fitness of two State One objects, we guarantee that a measure of trait-fitness covariance is transported up to the higher, pairwise level. By inhibiting the extent to which the two objects can compete in virtue of differential reproductive fitness, a fitness alignment mechanism forces natural selection to act, if at all, on the pair of objects as a whole.

Previous accounts have recognized the centrality of assortment in sculpting evolutionary outcomes (Hamilton 1971; Frank 1998; Fletcher and Doebeli 2006), in terms of its effect on the evolution of altruism. However, little is done to explain why altruism – a sort of cooperative interaction between agents – should be more or less relevant than any other properties we associate with higher-level organization, for example, physical cohesion, autonomy, goal directedness, etc. The advantage of invoking Price's/Lewontin's perspectives is that they allow us to *justify* our identification of mechanisms such as assortment as being fundamental to the first steps of a transition process. Assortment is relevant *because* it drives selection up to a higher level, by shifting the expression of heritable variance in fitness up to the between-aggregate level.

Prima facie there are two conceptually distinct modes by which the fitness of separate particles may be brought into alignment. One is to take two objects which already are close in fitness in virtue of common expression of some fitness-affecting and heritable trait, and to introduce a mechanism which inhibits them from subsequent fitness divergence. This is really a mere sorting mechanism, something which enables the objects to preferentially assort with objects that are similar to themselves. Kin recognition is an oft-discussed mechanism of this sort, but simpler factors such as substrate viscosity can also suffice to bring about fitness correlation just by making objects more likely to interact with, and share an environment with, close kin (Nowak and May 1992). Finally, ancestral kinship can be absent entirely but assortment take place on the basis of traits that are similar without owing their similarity to shared inheritance (Gardner and West 2010).¹⁵

The other, distinct, mode is to introduce a mechanism which itself brings about convergence between the fitnesses of two objects. In this case, the objects need not start out similar, in fitness terms or otherwise. What is essential is that they enter into some sort of causal relationship or 'fitness-affecting interaction' (Sober and Wilson 1998) which makes the evolutionary future of each depend on the other. Then a loss to one partner implies a loss for the other, and vice versa. This mode of fitness alignment will especially underpin the early aggregation of objects which are very different from one another, as occurs in symbioses (Bull and Rice 1991; Frank 1994; Foster and Wenseleers 2006). We can,

following Queller, call this mode of aggregation 'egalitarian', in contrast to the preceding 'fraternal' mode (Queller 2000). The egalitarian form is more subject to conflict, although there might be more ways to generate cooperative fitness benefits, just because the different parts need not be identical (Calcott and Sterelny 2011). However, extant higher-level organisms exemplify both modes of transition.

In real life we may expect that these two modes will be combined in actual interactions. Common to both types of mechanism is that they can provide a first step, a way to answer the how-question, because they will each qualify a population as having left State One, by delineating groups in which some heritable fitness variance will show up at the between-group level. They introduce the possibility of higher-level selection. Neither type of mechanism will actually be selected without a further, why-condition being met.¹⁶ It has to be true that some advantage is associated with the mechanism. It is with respect to the why-question that the issue of synergy often asserts itself. Transitional steps are selected when the fitness alignment device enables partners to enjoy some good that they could not have achieved alone. The fitness of the whole is then more than the sum of its parts, whether because they are cooperating, or using division of labour to make efficiency gains, or simply taking advantage of a nonlinear relationship between group size and some good (Michod 2006).

5. An evolutionary chicken and egg

It appears that a chicken and egg problem threatens attempts to explain evolutionary transitions in individuality (Sober and Wilson 1998; Gardner and Grafen 2009; Rainey and Kerr 2011; Trestman 2013). The problem is as follows.

A key insight that has been gained via reflection on the significance of major transitions for our general understanding of evolution is that we must not presuppose the existence of higher-level organisms when offering evolutionary explanations (Griesemer 2000; Okasha 2001). We are obliged to explain the appearance of higher-level organisms, rather than presuppose them as targets in explanations of how selection can act at a higher level. In order not to illegitimately presuppose the existence of aggregate, higher-level organisms, we need to be able to say what causes an evolutionary transition – what drives movement, over evolutionary time, from State One to State Two populations. One popular idea, arrived at especially via consideration of the third class of transition problematics (about how transitions are maintained), is that what explains the possibility of higher-level organisms is their possession of policing mechanisms, or mechanisms of fitness alignment. Bottleneck developmental stages, for example, can

¹⁵ Note that this still qualifies as kin selection under the statistical understanding of relatedness as a regression coefficient (Frank 1998).

¹⁶ Although a new mechanism may spread through a population just as a consequence of drift or novel niche occupation. I thank my referee for pointing this out.

answer the problem of how multicellular organisms remain stable in the face of lower-level free-rider problems, by constraining the extent to which the parts of multicellulars show heritable variance in fitness (Maynard Smith and Szathmáry 1995). Another example is the germ soma separation observed in higher vertebrates. By severing the heritability of somatic cells, somatic differentiation contains the evolutionary hazards that would otherwise be posed by selfish somatic mutants (Buss 1987). An example from a different level of organization is fair meiosis – a complex mechanism which guarantees that no gene gets a preferential fitness advantage during cell mitosis (Leigh 1971). In general, individuating mechanisms explain how transitioning organisms avoid being undermined by lower-level competition (Clarke 2013).

However, this solution contains the seed of a problem. For mechanisms such as bottlenecks and somatic differentiation also stand, themselves, in need of explanation. They are not properties that we may presuppose as having being built into life since its first inception. Rather, they are derived phenomena whose appearance on the evolutionary scene must itself be accounted for. On the face of it, this is not a special problem. Biologists are accustomed to explaining the emergence of complex phenotypic traits. Indeed, many see this as the primary triumph of evolutionary theory. Just as with other complex and beneficial traits, the natural response is to view developmental bottlenecks and germ lines as adaptations – traits which have emerged gradually as a consequence of the cumulative action of natural selection (Williams 1966). Note further that given their role in suppressing within-organism competition, we cannot treat bottlenecks or germ lines as having being selected by within-organism (between-cell¹⁷) selection. Instead they qualify naturally as higher-level adaptations – traits which have emerged as a consequence of selective pressures acting at the between-organism level.

The problem with this line of reasoning – offering bottlenecks up as an example of something which can drive selection up to the higher level, and then accounting for the existence of bottlenecks by appeal to the notion of adaptation – is that it presents bottlenecks as both causes and consequences of higher-level selection (Sober and Wilson 1998; Gardner and Grafen 2009). But each of these roles requires the prior occupation of the other. In other words, which came first, the higher-level selection or the bottlenecked organism?

In fact this apparent problem is not real. The appearance of a chicken and egg problem arises from the expectation that properties of very complex organisms must be explicable right at beginning stages of a transition process. For example, Michod's models assume that the properties of complex organisms – germ soma separation, division of

labour, collective reproduction – are present in the proto-multicellular organism (Michod 1999). Instead we need to realize that different properties will apply to higher-level individuals in the early stages of a transition process from those that apply later on, when the transition is complete, or nearly complete. Higher-level individuals do not spring into existence fully complex, with lower-level conflict suppressed a hundred per cent.

Complex mechanisms, such as bottlenecks and germ lines, as well as worker policing, and immune surveillance, will be present only in organisms with very high degrees of individuality – those that have made a large amount of progress along a transition continuum. They are derived properties, which should be treated as higher-level adaptations, that is, as consequences of higher-level selection. However, we go wrong in presupposing mechanisms such as these as necessary *causes* of higher-level selection. This role is played instead by very simple individuating mechanisms – things which are sufficient to cause the right sort of fitness alignment among a population of State One objects, but which can arise by chance alone, without having to be cumulatively selected for. These very simple mechanisms are causes, but not consequences, of higher-level selection.¹⁸ I give examples of such mechanisms in section 6.

Once the new selective level is operational, a possible consequence is the cumulative selection of traits which make the level robust to interference from lower-level cheats. There is the potential to get positive feedback, in which early-stage traits enable the selection of incrementally more complex traits (Crespi 2004). In other words, a self-fuelling process can initiate in which selection for higher-level advantages, such as predator escape, drives the eventual evolution of complex individuating mechanisms such as bottlenecks and germ lines.¹⁹ 'The coevolution of traits that influence population structure with traits that are favoured by the new population structure can result in a feedback process that concentrates natural selection at one level of the biological hierarchy' (Sober and Wilson 1998 p 97). At the end of a long self-ratcheting process we can expect to see very complex mechanisms, which could not possibly have occurred by chance, and yet which do not circularly act as both cause and effect of higher-level selection.

Possible outcomes include differentiated parts/division of labour, differentiated life stages/life cycle, reproductive specialization of parts, functional integration, synergy (the group output being a nonlinear function of members' output),

¹⁷ Again, I refer here to selection between cells *within the organism*, rather than globally (Sober 2011).

¹⁸ As my referee pointed out, these simple mechanisms may be exaptations – phenomena that were selected at the lower level for one function, but are subsequently recruited for a different purpose in the higher-level collective.

¹⁹ Such feedback processes have now been modelled, for example in García and De Monte 2013; Powers *et al.* 2011; Fleming 2012.

bottlenecked propagule reproduction and other paradigmatically ‘emergent’ properties. Such properties need not be simply more of the same, with respect to the early-stage benefits, but can be qualitatively different in character. For example, Birch carves out an important practical role for what he terms ‘extreme redundancy’—where there are more workers than are strictly needed for task completion—in transforming collectives from groups into individuals. Redundancy generates synergy by increasing the robustness of task completion (Birch 2012). Similarly, divisions of labour and specialization, task-based cooperation (Anderson *et al.* 2001) and increases in evolvability (the extent to which a lineage is able to explore new phenotypes) are all advantageous states of affairs that are possible once objects have begun to express fitness variance between groups rather than within them.

There is no chicken and egg complication as long as we conceptualize organismality as something which emerges gradually out of reciprocal feedback between higher-level selection and the mechanisms which support it. Complex, late stage organismal traits may well be products of higher-level selection, but are not prerequisites for selection at the higher level. Simple early stage traits such as stickiness, on the other hand, are prerequisites of higher-level selection, but they need not be products of it.

6. Transitioning: A beginner’s guide

Now it is time to add some more concrete details to the abstract schema I have been building. My aim is to describe some plausible steps by which an actual population could begin to transition. As explained, I need to describe a simple novel trait which could begin to align the fitnesses of some State One objects. The mechanism may either assort things that already have similar fitness or it can cause their fitnesses to converge. Crucially, it must be the sort of trait that could arise by chance or mutation (or some sort of self-organization such as a phase-transition (Newman and Bhat 2011; Nanjundiah 2013)), without demanding a cumulative selection process.

My first how-possibly mechanism is stickiness. Imagine that one cell in a freely mixing population acquires a mutation that causes its bearer to become slightly sticky. Perhaps it produces an adhesive molecule, a glue. Next time the cell divides, its daughter might remain stuck to it, instead of separating. Alternatively, instead of binding mitotic progeny, the glue could bind unrelated cells. In both cases, a new aggregate two-cell object has been created. Further mitotic division events might cause the aggregate to grow in cell number indefinitely, because the mitotic daughters of the original mutant inherit the stickiness. At some point, depending on the adhesivity of the glue, the aggregate is likely to fission into separate pieces. Assuming that the stickiness is compatible with the cells achieving the minimal

requirements for viability and fecundity, the separate pieces will each reiterate a similar trajectory, growing and then fissioning, although the precise rates will depend upon contingencies of how/where the fission occurred, and on just how sticky each of the cells in the aggregates is. The cells in this population are no longer freely mixing—the stickiness has brought about some population structure.

Much attention has already been given to the notion of incomplete separation after mitosis, in constituting a first step, especially in the transition to multicellularity, but why does stickiness qualify as a simple individuating mechanism? In other words, what is it about stickiness that makes it a first step along a transition continuum—the first step of a change in the hierarchical level of selection?²⁰ The answer is that it gives rise to fitness alignment. The glue may bind cells that already share fitness in virtue of common ancestry (fraternal mode), but even in the absence of kinship, the glue aligns the fitness of the cells it binds by forcing them to share an environment (egalitarian mode). Two cells that remain physically attached for the duration of their cellular life are more likely to experience similar environments, and so to manifest a similar lifetime rate of reproduction, than are two cells that are free to drift apart. Good things that happen to one cell are highly likely to also happen to the other cell. The stickiness, or rather the glue, therefore constitutes a mechanism which couples their fates, to some small extent. Their fitnesses can still diverge—one cell may acquire a mutation that enables it to reproduce faster, and fail to share it with its neighbour. Nonetheless, the possibilities for divergent evolutionary fates are smaller, in the presence of the glue, than in a situation where the mutant can float away and experience different conditions from its partner. The glue brings about some small measure of correlation, so that a cell is statistically more likely to have a fitness that is similar to the other cell in its aggregate than it is to another random cell from the population. To put it another way, we would measure, on average, more fitness variance if we compare cells in different aggregates than if we compare cells within aggregates.

Assortment alone does not guarantee that the novel mechanism will spread through the population. For selection to favour the glue, it needs to be the case that there is an advantage—that sticky aggregates do better, for some reason, than non-sticky individual cells. This is where class two questions come in—what is the evolutionary benefit? Possible candidate advantages include: avoiding predation because a group of cells may not fit into the mouth of the

²⁰ I deny that stickiness, on its own, can bestow one hundred per cent higher-level individuality on an aggregate of cells (contrast Newman 2003). It is merely a first step. My argument is that a whole transition—all the way to the capacity for selection being held exclusively at the higher level—requires a multiplicity of such steps, and in this sense should be viewed as gradual, rather than happening all at once.

cells' usual predators; enhanced motility, dispersal or weight allowing easier access to a preferred environment (Ratcliff *et al.* 2012); enhanced storage/absorption of some nutrient (Michod 2006); slower heat loss. It must be stressed, however, that whatever advantage is gained, it must outweigh any disadvantages, of which a similar list can be generated: decreased surface area; faster cumulative build-up of waste products; greater demand for resources; facilitated predation and parasitization via localization. Sticking together will often be a bad idea! For example, if a glue-producing wrinkly spreader type of *Pseudomonas* cell grows under shaken conditions, it gets sick (Rainey and Kerr 2011). Stickiness will only spread through the population if it bestows some net advantage on the cells that express it. But equally, a benefit is of no consequence at all in the absence of an assortment mechanism. When both come together, then there is scope for higher-level selection to increase the frequency of the aggregates. Once a small amount of higher-level selection obtains, a new possibility emerges. One of the things that might get selected at the new higher level is a mechanism that makes the aggregation more robust.

So stickiness provides an example of a simple, early-stage individuating mechanism – something that will align the fitness of objects in a population. Once a glue-producing mutation has occurred a first step has been taken, because its bearer now has the properties necessary for generating a selectable higher-level population.

Stickiness may be sufficient in allowing a population to set off on a transition trajectory, but it is not essential, because many other mechanisms could play the same role. A commitment to future fitness-affecting interactions can be wrought by other mechanisms of physical constraint, such as the imposition of common boundaries. A physical barrier achieves the requisite effect on population-level variance by removing the possibility of migration between groups after they have formed (Leigh 1971). The appearance of the cell wall may thus have been a key event in the transition from pre-cellular to cellular organization. Engulfment by another object is a fitness alignment mechanism that has driven endosymbiotic transition events (Margulis 1967). Another, even simpler, cause of spatial constraint is substrate viscosity (Taylor 1992). If objects are dividing on or in a medium that allows only limited dispersal, then kinship can ensure that a correlation of types emerges, such that it can underpin the higher-level selection of advantageous novel traits such as production of common goods, and other cooperative acts. Cells do not need to be physically attached to one another in order to form selectable siderophore-producing groups, for example (Griffin *et al.* 2004).

Furthermore, committed fitness-affecting interactions can emerge in the absence of anything as crude as physical glues or obstacles. In social insect colonies virtual

boundaries are maintained by complex hormonal and behavioural systems which bring about inter-colony aggression (Tsutsui *et al.* 2003). Higher vertebrates police their boundaries and monitor interactions among their various parts using an incredibly complex innate immune system (Pradeu 2010). Another way to couple the fitnesses of two, initially different, objects is to couple their reproduction. Parasites that are obligately vertically descending with their host, for example, acquire a common fate with that host. This is a sort of fitness-aligning interaction, even if the host fitness is entirely unchanged (Frank 1997). In metazoans with sufficient cognitive complexity, social bonds can be achieved via direct reciprocity (such as grooming behaviours in monkeys) or indirect reciprocity (e.g. reputation effects; Nowak and Sigmund 1998).

The common effect shared by all these mechanisms is fitness alignment. The device in each case guarantees that there is some limit on the extent to which the parts of the group can diverge in fitness. To achieve a level-general characterization of an evolutionary transition's first step we must remember that different concrete mechanisms are likely to be used to align fitness in different lineages, and at different levels. The transition to eusociality in insects cannot be explained in terms of glue. Termites, wasps and ants do not thrive in colonies by sticking physically to the body of their mother. Instead their fitnesses are aligned in virtue of ongoing behavioural interactions. However, interesting distinctions can be drawn between different mechanisms, in terms of their differential consequences for heritability and thus for subsequent evolution. For example a complex fitness-alignment mechanism such as a single-celled bottleneck implies different constraints on subsequent evolution than does a simple mechanism such as stickiness. Bottlenecks can be effective in enhancing the heritability of traits across generations, by limiting the extent to which heritability is undermined by transmission bias (Frank 2012a). If we are considering a population, on the other hand, in which there is occasional dispersal of glue-producers into a global pool, followed by re-assortment, then higher-level heritability is likely to be very low, unless there is some further mechanism for aligning like with like at the assortment stage. An aggregate that divides by random fission allows much less higher-level heritability than a more organized mechanism of multiplication might, which makes fragmentation periodic, or brings it under endogenous control. So the mechanisms are far from equivalent in terms of successive evolutionary outcomes.

7. A note on fitness and reproduction

My model in section 3 permits use of an MSL1 definition of higher-level fitness (Damuth and Heisler 1988) where

the higher-level rate of reproduction is understood as the average of the rates of reproduction by member objects. Models of MLS1-type selection consider differential growth of collectives as sufficient for generating change in trait frequencies over time. In addition to the concerns over ‘cross-level by-products’, which I have already addressed, some people complain that MLS1 analysis fails to delineate ‘real’ higher-level selection, because there is no emergence of true higher-level reproduction (‘export of fitness’) (Michod 1999; 2005; Simon *et al.* 2012; Libby and Rainey 2013). In effect, such people demand that we declare evolutionary change as being caused by higher-level selection only at the stage when the higher-level objects are behaving truly analogously to the objects in State One. My problem with this sort of view is that it asks the impossible. It essentially expects a population to be able to move directly from State One to State Two, without passing through intermediate states. State Two might well be the eventual destination of a transitioning population, but we cannot reasonably demand that the population meets State Two conditions *as a first step*. We do not want, after all, to presuppose highly derived properties as pre-existing conditions of early higher-level organisms. One advantage of focusing on early stages is that it allows us to avoid difficult questions about what it means for a property or process to be genuinely ‘emergent’. Between a fleet herd of deer and a herd of fleet deer (Williams 1966) are many grades of ambiguity. The concept ‘reproduction’ carries many intuitive connotations that are carried from the highly derived life cycle of complex metazoans, and that are best avoided. For these reasons, I do not make MLS2-style reproduction a necessary capacity of first-step aggregates.

Note, however, a point that is not usually given sufficient attention in MLS1 models: If selection is purely a process of differential growth/longevity of MLS1 aggregates,²¹ then there will be a clear finite limit to the process. No further evolution will be possible beyond the point at which the population contains only one aggregate. Variance is not definable, after all, for a population of one. Similarly, no Darwinian evolution can occur in a population where aggregate-formation is always followed by the death of all the objects in the aggregates, as occurs in Libby and Rainey’s model of mat formation in *Pseudomonas* populations (Libby and Rainey 2013). There is no possibility of selection for aggregate-level traits if aggregates are always evolutionary deadends. For this reason, we have to build in the condition that there will be some mechanism for maintaining a multiplicity of aggregates.

Aggregate reproduction, understood in a paradigm MLS2 sense, is obviously a sufficient mechanism for meeting this

condition. Nonetheless, a much more minimal mechanism will do. In the case of the ‘stickiness’ mechanism, it is enough to assume that the aggregates will inevitably fission at some point, and this might plausibly be as a mere consequence of the generation of unsustainable shear forces once the aggregate exceeds a certain density of lower-level parts. Alternatively, it could be that the glue is imperfect, so that offspring State One objects sometimes get stuck to parent and sometimes do not. Or perhaps the strength of the glue depends on some environmental factor, so that groups occasionally break up and reform later. Some people may want to call even these very rudimentary process ‘reproduction’, but the important point is that it need not be in any sense (endogenously) developmentally controlled, or regular/periodical, or such that clear parent-offspring distinctions can always be drawn. There is no need to specify anything about whether fission should produce equal-sized-parts, or be more a matter of small pieces detaching from the parent. Such minimal fragmentary processes can be very far away from what we might call ‘paradigm’ or ‘full-blown’ reproductive events.²² Furthermore, higher-level multiplicative capacity, in this very minimal sense, does not require any additional steps/mutations. It comes ‘for free’ as a consequence of the reproductive capacities of the lower-level objects, given some very undemanding assumptions about shear forces or environmental contingencies.²³ Group multiplicity may just be an inevitable consequence of cell multiplicity, given enough time.

8. Conclusions

I have sketched out the ideal anatomy for an explanation of an evolutionary transition in individuality, during which solitary organisms become subsumed within a new higher-level organism, which participates in a higher-level selection process. Any successful account of an evolutionary transition must model it as a process which happens gradually, proceeding from very simple first steps. I defined primitive/first-step higher-level organismality in terms of the minimal conditions necessary for participating in a higher-level selection process. The very first step consists in the appearance of some trait which marks the difference between a population in which selection acts only at the lower level, and one in which selection (also) acts at a higher level.

I explained how these definitions allow us to overcome a conceptual ‘chicken and egg’ problem which arises from the

²¹ Bouchard 2008.

²² In any case, even paradigm reproduction would be insufficient in the case where offspring production is exactly coincident with the death of the parent, so that the replication event fails to make any positive increase in the overall population number.

²³ Coral colonies, for example, do not fission as a programmed developmental event, like starfish do. It just occurs when the group is large and there is a storm.

expectation that properties of very complex organisms must be explicable right at beginning stages of a transition process. Instead we need to realize that different properties will apply to higher-level individuals in the early stages of a transition process from those that apply later on, when the transition is complete, or nearly complete.

Acknowledgements

This article emerged out of long and fruitful discussions with Silvia De Monte and Paul B Rainey, which followed the hugely fecund meeting of perspectives that was ‘Individuals and Groups’ in Almora. In addition to them, I give many thanks to Stuart Newman, my anonymous referees and especially to Vidyanand Nanjundiah.

References

- Anderson C, Franks NR and McShea DW 2001 The complexity and hierarchical structure of tasks in insect societies. *Anim. Behav.* **62** 643–651
- Bijma P and Aanen DK 2009 Assortment, Hamilton’s rule and multilevel selection. *Proc. R. Soc. B.* **277** 673–675
- Birch J 2012 Collective action in the fraternal transitions. *Biol. Philos.* **27** 363–380
- Bonner JT 1974 *On development: The biology of form* (Harvard, MA: Harvard University Press)
- Bonner JT 1998 The origins of multicellularity. *Integr. Biol.* **1** 27–36
- Bouchard F 2008 Causal processes, fitness and the differential persistence of lineages. *Philos. Sci.* **75** 560–570
- Bourke AFG 2011 *Principles of social evolution* (Oxford: Oxford University Press)
- Bull JJ and Rice WR 1991 Distinguishing mechanisms for the evolution of co-operation. *J. Theor. Biol.* **149** 63–74
- Buss L 1987 *The evolution of individuality* (Princeton, NJ: Princeton University Press)
- Calcott B 2008 The *other* cooperation problem: Generating benefit. *Biol. Philos.* **23** 179–203
- Calcott B 2011 Alternative Patterns of explanation for major transitions; in *The major transitions in evolution revisited* (eds) B Calcott and K Sterelny (Cambridge, Massachusetts: The MIT Press)
- Calcott B and Sterelny K 2011 Introduction: A dynamic view of evolution; in *The major transitions in evolution revisited* (eds) B Calcott and K Sterelny (Cambridge, Massachusetts: The MIT Press) pp 1–14
- Clarke E 2013 The multiple realizability of biological individuals. *J. Phil.* **8** 413–435
- Crespi BJ 2004 Vicious circles: positive feedback in major evolutionary and ecological transitions. *TREE* **19** 627–633
- Damuth J and Heisler L 1988 Alternative formulations of multilevel selection. *Biol. Philos.* **3** 407–430
- Dawkins R 1982 *The extended phenotype* (Oxford: Oxford University Press)
- Fleming L 2012 Network Theory and the formation of groups without evolutionary forces. *Evol. Biol.* **39** 94–105
- Fletcher JA and Doebeli M 2006 How altruism evolves: assortment and synergy. *J. Evol. Biol.* **19** 1389–1393
- Fletcher JA and Doebeli M 2009 A simple and general explanation for the evolution of altruism. *Proc. R. Soc. B* **276** 13–19
- Fletcher JA and Doebeli M 2010 Assortment is a more fundamental explanation for the evolution of altruism than inclusive fitness or multilevel selection: reply to Bijma and Aanen. *Proc. R. Soc. B* **277** 677–678
- Folse HJ III and Roughgarden J 2010 What is an individual organism? A multilevel selection perspective. *Q. Rev. Biol.* **85** 447–472
- Foster KR 2009 A defense of sociobiology. *Cold Spring Harbor Symp. Quant. Biol.* **74** 403–418
- Foster KR and Wenseleers T 2006 A general model for the evolution of mutualisms. *J. Evol. Biol.* **19** 1283–1293
- Frank SA 1994 Genetics of mutualism - the evolution of altruism between species. *J. Theor. Biol.* **170** 393–400
- Frank SA 1997 Models of symbiosis. *Am. Nat.* **150** 80–99
- Frank SA 1998 *The foundations of social evolution* (Princeton, New Jersey: Princeton University Press)
- Frank SA 2012a Natural Selection. IV. The Price Equation. *J. Evol. Biol.* **25** 1002–1019
- Frank SA 2012b Natural Selection. III. Selection versus transmission and the levels of selection. *J. Evol. Biol.* **25** 227–243
- Garcia T and De Monte S 2013 Group formation and the evolution of sociality. *Evolution* **67** 131–141
- Gardner A 2008 The price equation. *Curr. Biol.* **18** R198–R202
- Gardner A and Grafen A 2009 Capturing the superorganism: a formal theory of group adaptation. *J. Evol. Biol.* **22** 659–671
- Gardner A and West S 2010 Greenbeards. *Evolution* **45** 25–38
- Gardner A, West SA and Wild G 2011 The genetical theory of kin selection. *J. Evol. Biol.* **24** 1020–1043
- Gilbert S, Sapp J and Tauber A 2012 A symbiotic view of life: we have never been individuals. *Q. Rev. Biol.* **87** 325–341
- Godfrey-Smith P 2008 Varieties of Population Structure and the Levels of Selection. *Brit. J. Phil. Sci.* **59** 25–50
- Godfrey-Smith P 2009 *Darwinian populations and natural selection* (Oxford: Oxford University Press)
- Goodnight C 2013 On multilevel selection and kin selection: Contextual analysis meets direct fitness. *Evolution* **67** 1539–1548
- Goodnight CJ, Schwartz JM and Stevens L 1992 Contextual analysis of models of group selection, soft selection, hard selection and the evolution of altruism. *Am. Nat.* **140** 743–761
- Griesemer J 2000 The units of evolutionary transition. *Selection* **1** 67–80
- Griffin AS, West SA and Buckling A 2004 Cooperation and competition in pathogenic bacteria. *Nature* **430** 1024–1027
- Grosberg RK and Strathmann RR 2007 The evolution of multicellularity: a minor major transition? *Annu. Rev. Ecol. Evol. Syst.* **38** 621–654
- Hamilton WD 1964 The genetical evolution of social behaviour, I. *J. Theor. Biol.* **7** 1–52
- Hamilton WD 1971 Selection of selfish and altruistic behavior in some extreme models; in *Man and beast: Comparative social behaviour* (eds) JF Eisenberg and WS Dillon (Washington, DC: Smithsonian Press)

- Hamilton A, Smith NR and Haber MH 2009 Social insects and the individuality thesis: Cohesion and the colony as a selectable individual; in *Organization of insect societies: From genome to sociocomplexity* (eds) J Gadau and F Fewell (Harvard USA: Harvard University Press) 572–589
- Heisler L and Damuth J 1987 A method for analyzing selection in hierarchically structured populations. *Am. Nat.* **130** 582–602
- Kerr B and Godfrey-Smith P 2002 Individualist and multi-level perspectives on selection in structured populations. *Biol. Philos.* **17** 477–517
- Kirk DL 1998 *Volvox: The molecular genetic origins of multicellularity and cellular differentiation* (Cambridge: Cambridge University Press)
- Kirk DL 2003 Seeking the ultimate and proximate causes of Volvox multicellularity and cellular differentiation. *Integr. Comp. Biol.* **43** 247
- Kirk DL 2005 A twelve-step program for evolving multicellularity and a division of labor. *Bioessays* **27** 299–310
- Kirk DL and Harper JF 1986 Genetic, biochemical and molecular approaches to *Volvox* development and evolution. *Int. Rev. Cytol.* **99** 217–293
- Koufopanou V 1994 The evolution of soma in the volvocales *Am. Nat.* **143** 907–931
- Larson A, Kirk MM and Kirk DL 1992 Molecular phylogeny of the volvocine flagellates. *Mol. Biol. Evol.* **9** 85–105
- Lehmann L, Keller L, West S and Roze D 2007 Group selection and kin selection: Two concepts but one process. *PNAS* **104** 6736–6739
- Leigh EG Jr 1971 *Adaptation and diversity* (San Francisco, CA: Freeman Cooper)
- Leimar O and Hammerstein P 2010 Cooperation for direct fitness benefits. *Phil. Trans. R. Soc. B* **365** 2619–2626
- Lewontin RC 1970 The units of selection. *Annu. Rev. Ecol. Syst.* **1** 1–18
- Libby E and Rainey P 2013 A conceptual framework for the evolutionary origins of multicellularity. *Phys. Biol.* **10** 035001
- Margulis L 1967 On the origin of mitosing cells. *J. Theor. Biol.* **14** 255–274
- Marshall JAR 2011 Group selection and kin selection: formally equivalent approaches. *TREE* **26** 325–332
- Maynard Smith J and Szathmáry E 1995 *The major transitions in evolution* (Oxford: Oxford University Press)
- Mayr E 1982 *The growth of biological thought* (Cambridge USA: Balknap Harvard)
- Michod RE 1999 *Darwinian dynamics: evolutionary transitions in fitness and individuality* (Princeton, USA: Princeton University Press)
- Michod R 2005 On the transfer of fitness from the cell to the multicellular organism. *Biol. Philos.* **20** 967–987
- Michod RE 2006 Group covariance effect and tradeoffs during evolutionary transitions in individuality. *Nat. Acad. Sci.* **103** 9113–9117
- Nanjundiah V 2013 Review of ‘Transformations of Lamarckism. From Subtle Fluids to Molecular Biology’. *Curr. Sci.* **104** 523–525
- Newman SA 2003 Hierarchy; in *Keywords and concepts in evolutionary developmental biology* (eds) BK Hall and WM Olson (Cambridge MA: Harvard University Press) pp 169–174
- Newman SA and Bhat R 2011 Lamarck’s dangerous idea; in *Transformations of Lamarckism. From subtle fluids to molecular biology* (eds) SB Gissis and E Jablonka (Cambridge MA: The MIT Press)
- Nowak MA 2006 Five rules for the evolution of cooperation. *Science* **314** 1560–1563
- Nowak MA and Highfield R 2011 *Supercooperators: Evolution, altruism and human behaviour* (Edinburgh: Canongate Books Ltd)
- Nowak MA and May RM 1992 Evolutionary games and spatial chaos. *Nature* **359** 826–828
- Nowak MA and Sigmund K 1998 Evolution of indirect reciprocity by image scoring. *Nature* **393** 573–577
- Nunney L 1985 Group selection, altruism, and structured-deme models. *Am. Nat.* **126** 212–230
- Okasha S 2001 Why won’t the group selection controversy go away? *Brit. J. Philos. Sci.* **52** 25–50
- Okasha S 2005 Multil-level Selection and the Major Transitions in Evolution. *Philos. Sci.* **72** 1013–1025
- Okasha S 2006 *Evolution and the levels of selection* (Oxford: Oxford University Press)
- Powers ST, Penn AS and Watson RA 2011 The concurrent evolution of cooperation and the population structures that support it. *Evolution* **65** 1527–1543
- Pradeu T 2010 What is an organism? An immunological answer. *Hist. Phil. Life Sci.* **32**:247–268
- Price GR 1970 Selection and covariance. *Nature* **227** 520–521
- Price GR 1972 Extension of covariance selection mathematics. *Ann. Hum. Genet.* **35** 455–458
- Queller DC 2000 Relatedness and the fraternal major transitions. *Phil. Trans. Roy. Soc. B. Biol. Sci.* **355** 1647–1655
- Queller DC and Strassmann JE 2009 Beyond society: the evolution of organismality. *Phil. Trans. Roy. Soc. B. Biol. Sci.* **364** 3143–3155
- Rainey PB and Kerr B 2011 Conflicts among levels of selection as fuel for the evolution of individuality; in *The major transitions in evolution revisited* (eds) B Calcott and K Sterelny (Cambridge, Massachusetts: The MIT Press)
- Rainey PB and Rainey K 2003 Evolution of cooperation and conflict in experimental bacterial populations. *Nature* **425** 72–74
- Ratcliff WC, Denison RF, Borrello M and Travisano M 2012 Experimental evolution of multicellularity. *PNAS* **109** 1595–1600
- Ratnieks FLW 1988 Reproductive harmony via mutual policing by workers in eusocial Hymenoptera. *Am. Nat.* **132** 217–236
- Ratnieks FLW, Foster K and Wenseleers T 2006 Conflict resolution in insect societies. *Annu. Rev. Entomol.* **51** 581–608
- Rokas A 2008 The origins of multicellularity and the early history of the genetic toolkit for animal development. *Annu. Rev. Genet.* **42** 235–251
- Sachs JL 2008 Resolving the first steps to multicellularity. *TREE* **23** 245–248
- Sachs JL, Mueller UG, Wilcox TP and Bull JJ 2004 The evolution of cooperation. *Q. Rev. Biol.* **79** 135–160
- Simon B, Fletcher KA and Doebeli M 2012 Towards a general theory of group selection. *Evolution* **67** 1561–1572
- Sober E 2011 Realism, Conventionalism, and causal decomposition in units of selection: Reflections on Samir Okasha’s *Evolution*

- and the Levels of Selection. *Philos. Phenomenol. Res.* **82** 221–231
- Sober E and Wilson DS 1998 *Unto others: The evolution and psychology of unselfish behaviour* (Cambridge, MA: Harvard University Press)
- Taylor PD 1992 Altruism in viscous populations – an inclusive fitness model. *Evol. Ecol.* **6** 352–356
- Trestman M 2013 Which comes first in evolutionary transitions: The behavioural chicken, or the evolutionary egg? *Biol. Theory* **7** 48–55
- Tsutsui ND, Suarez AV and Grosberg RK 2003 Genetic diversity, asymmetrical aggression, and recognition in a widespread invasive species. *PNAS* **100** 1078–1083
- Van Veelen M, Sabelis MW and Egas M 2012 Group selection and inclusive fitness are not equivalent. *J. Theor. Biol.* **299** 64–80
- West SA, Griffin AS and Gardner A 2007 Evolutionary explanations for cooperation. *Curr. Biol.* **17** R661–R672
- Williams GC 1966 *Adaptation and natural selection: a critique of some current evolutionary thought* (Princeton: Princeton University Press)