OPEN ACCESS

University of BRISTOL

Okasha, S. (2012). Emergence, hierarchy and top-down causation in evolutionary biology. Interface Focus, 2(1), 49-54. 10.1098/rsfs.2011.0046

Link to published version (if available):
10.1098/rsfs.2011.0046

Link to publication record in Explore Bristol Research
PDF-document

University of Bristol - Explore Bristol Research

General rights

Title: Emergence, Hierarchy, and Top-Down Causation in Evolutionary Biology

Author: Samir Okasha, Department of Philosophy, University of Bristol, 9 Woodland Road, Bristol BS8 1TB, U. K.

Summary

The concept of emergence and the related notion of 'downward causation' have arisen in numerous branches of science, and have also been extensively discussed in philosophy. Here I examine emergence and downward causation in relation to evolutionary biology. I focus on the old but ongoing discussion in evolutionary biology over the 'levels of selection' question, which asks which level(s) of the biological hierarchy natural selection acts at, e.g. the gene, individual, group or species level? The concept of emergence has arisen in the levels of selection literature as a putative way of distinguishing between 'true' selection at a higher level from cases where selection acts solely at the lower level but has effects that perlocate up the biological hierarchy, generating the *appearance* of higher-level selection. At first blush this problem seems to share a common structure with debates about emergence in other areas, but closer examination shows that it turns on issues that are *sui generis* to biology.

Key index words

Emergence, Hierarchical Organization, Evolutionary Biology, Levels of Selection, Causation, Price's equation

Introduction

The concept of emergence has featured in many branches of science including physics, complex systems, and neuroscience. There is also a large philosophical literature on the topic [1]. I want to discuss emergence, and the closely related topics of hierarchical structure and 'top-down' causation, in relation to evolutionary biology. My concern is not with the whole of evolutionary biology, but rather with one particular debate within it – the on-going discussion over 'levels of selection' [2, 3, 4]. I start by offering a brief characterization of the

levels-of-selection question in biology. I then outline a central conceptual problem within the levels-of-selection debate, and explore how the notions of emergent property and top-down causation have arisen in the biological literature as possible solutions to the problem.

My central claim is this. At first blush, the issue within evolutionary biology that I discuss *appears* to be a special case of the more general issue of emergence/top-down causation, examined in the abstract by authors such as Jaegwon Kim [5]. However, closer examination reveals that this is not actually the case. In fact, the biological issue is largely *sui generis*, and turns on matters that are specific to evolutionary science, despite the appearance of a common structure with debates in other disciplines. This claim, if correct, dovetails nicely with an argument made in a different context by Larry Sklar, about the relation between philosophical issues as they arise within science and as they are treated in the abstract [6].

1. The Levels of Selection Question

The levels-of-selection question asks which level or levels of the biological hierarchy natural selection acts at. The question is a fundamental one for evolutionary biology, for it arises directly from the underlying logic of Darwinism. As Darwin himself realised, the principle of natural selection is entirely abstract: it tells us that if a population of 'entities' exhibits variation, differential reproduction and heredity, then its composition will change over time, as the 'fittest' variants gradually replace the less fit. In most discussions, including Darwin's own, these 'entities' are taken to be individual organisms, but in theory at least, there are other possibilities. For the biological world is hierarchically organised, with organisms somewhere in the middle. Each organism is composed of organs and tissues, which are themselves made up of cells; each cell contains a number of organelles and a cell nucleus; each nucleus contains a number of chromosomes; and on each chromosome lie a number of genes. Above the level of the organism we find entities such as kin groups, colonies, demes, species and whole ecosystems.

How exactly the biological hierarchy should be characterized, that is, which levels should be recognised and why, is a non-trivial issue. But one point is clear from the outset. Entities at various hierarchical levels, above and below that of the organism, can satisfy the conditions required for evolution by natural selection. For just as organisms gives rise to

other organisms by reproduction, so cells give rise to other cells by cell division, genes to other genes by DNA replication, groups to other groups by fission (among other ways), species to other species by speciation, and so on. Thus the Darwinian concept of fitness, i.e. expected number of offspring, applies to entities of each of these types. So in principle, these entities could form populations that evolve by natural selection.

Historically, the levels-of-selection question has been closely linked with the problem of altruism. In biology, altruism refers to behaviours which reduce the fitness of the organism performing them but boost the fitness of others, e.g. sharing food. Selection at the level of the individual organism should disfavour altruism, for altruists suffer a fitness disadvantage relative to their selfish counterparts, yet such behaviour is quite common in nature. One possible explanation, first canvassed by Darwin himself, is that altruism may have evolved by selection at higher levels of organization, for example, the group or colony level. Groups containing a high proportion of altruists might have a selective advantage over groups containing mostly selfish types, thus allowing altruism to prosper. The idea that group selection might explain the evolution of altruism is still discussed today.

For many years, the idea of selection operating at levels other than that of the individual organism was seen as a theoretical curiosity, unlikely to be important in practice. Recently there has been a significant change of opinion among (some) biologists, and a resurgence of interest in hierarchical or 'multi-level' approaches to natural selection. This is for two main reasons. Firstly, theorists concerned with explaining the 'major evolutionary transitions' have realised that selection acting at multiple hierarchical levels (multi-level selection) may have played a major role [3, 4, 7, 8, 9]. Such transitions occur when a number of lower-level units, capable of surviving and reproducing alone, aggregate into a single larger unit, which eventually becomes a new higher-level individual. Many such transitions have occurred in the history of life (e.g. from single-celled to multi-celled organisms), giving rise to the modern biological hierarchy. For an evolutionary transition to occur, it is generally necessary for selection at the higher level to 'trump' selection at the lower level, to ensure that the lower-level units work for the good of the whole. From this perspective, we see that multi-level selection, far from being a theoretical curiosity, is in fact implicated in some of the most important evolutionary events on earth.

The second reason behind the resurgence of multi-level selection stems from a theoretical advance made by George Price in the 1970s, whose full significance has only recently been fully appreciated [10]. Price showed how the overall evolutionary change, in a population with hierarchical structure, could be partitioned into a number of components, one corresponding to each level of the hierarchy. Thus for example, if the two levels are individuals and groups, then applying Price's technique allows us to express the total evolutionary change as the sum of two components, one reflecting selection acting on individuals within groups, the other reflecting selection acting between the groups themselves. (For an 'altruistic' trait these two components will be opposite in sign, i.e. group selection will favour the trait, individual selection will oppose it.) As well as being a useful modelling technique, Price's analysis also provides a key conceptual insight: natural selection, at any hierarchical level, requires a *covariance between some trait and the fitness of entities at that level*. Moreover, the component of the total change due to selection at any level is directly proportional to the magnitude of the trait-fitness covariance at that level. I have provided a fuller account of Price's analysis of multi-level selection in previous work [2], as have other authors [3].

2. Emergence and the Levels of Selection

What has all this got to do with emergence and top-down causation? Interestingly, these notions have arisen in both the biological and the philosophical literature on levels of selection. In some ways this is not surprising, since the concept of emergence is potentially applicable wherever there is hierarchical structure, and as we have seen, the hierarchical nature of the biotic world is part of what gives rise to the levels of selection question. Moreover, causation is also central to the levels question. Darwinian explanations are usually understood as causal: to attribute the spread of a trait to natural selection is to say what *caused* it to spread. (This causal dimension to Darwinian explanations is not always made explicit, but it is generally intended.) If this is right, then in a multi-level scenario, where natural selection is operating at two (or more) levels of the biological hierarchy, it follows that two distinct causal processes are occurring, at different hierarchical levels; in principle, the higher-level process could impinge on entities at the lower level. And this is precisely the sort of situation for which the notion of top-down causation was tailored. So it is easy to see,

in general terms, why emergence and top-down causation should feature in discussions of the levels of selection in biology.

To make this more concrete, consider a well-known distinction introduced by G.C. Williams in his famous book *Adaptation and Natural Selection* [11]. Williams argued that group adaptations must be sharply distinguished from what he called 'fortuitous group benefits'. A group adaptation is a feature of a group that benefits it, and that evolved by selection at the group level. For example, some insect colonies have a sophisticated division-of-labour among workers, which probably evolved because of the advantage it confers on the whole colony; if so, this is a group adaptation. A fortuitous group benefit, by contrast, is a feature of a group that benefits it, but is not the result of group-level selection, rather it is an 'unintended side effect' of some other process. For example, if a particular deer herd contains deer than can run especially fast, then the average running speed of the herd will be high – higher than that of other herds. But 'average running speed' is not an adaptation of the deer herd. Rather, running fast is an adaptation of the individual deer within the herd; there is an individual-level selection story to be told about why they evolved to run fast. The fact that the herd they live in has a high average running speed is simply a side effect of the adaptations of the individual deer. No group-level selection process need be invoked to explain this fact.

One natural thought is that Williams' distinction between group adaptation and fortuitous group benefit lines up with the distinction between emergent and non-emergent (or 'aggregate') properties; this has been argued by a number of authors including Elizabeth Vrba [12]. Division-of-labour among the workers in an insect colony is arguably an emergent property of the whole colony, for it 'emerges' from the different activities and morphologies of many insects. But the average running speed of a deer herd is not like this – rather, it is a mere statistical aggregate of the running speeds of the individual deer. How exactly this emergent / aggregate distinction should be drawn, in biology and elsewhere, is a rather tricky question – for in both cases, the group property is presumably *determined* by individual properties. But if we grant that the distinction makes sense, and that we have at least a rough idea of how to apply it, then the hypothesis that emergent/aggregate coincides with the group adaptation/fortuitous group benefit distinction can at least be entertained. The hypothesis is *prima facie* quite plausible, for convincing examples of biological adaptations, at any level, are usually complex traits – and a hall-mark of a complex trait is precisely that it is not a

simple aggregation of lower-level traits, in the way that the herd's average running speed is a simple aggregation of the traits of the individual deer.

3. Causation and Cross-level Byproducts

Closely related to the emergent property issue is the issue of causation as it relates to levels of selection. To focus the issue, consider a two-level scenario, where lower-level 'particles' are nested within higher-level 'collectives', as in Figure 1 below. This figure could represent individual organisms within social groups, or cells within multi-celled organisms, or genes within genomes, or species within ecosystems. In principle, natural selection could operate on particles within collectives, or on whole collectives, or at both levels simultaneously. Price's analysis teaches us that for selection to act at a level, there must a trait-fitness covariance at that level. So for collective-level selection to occur, it is necessary that the fitness of a collective depend systematically on the traits of that collective. (Note that the 'traits' of a collective may include both aggregate properties – such as the proportion of particles of a given type it contains – or emergent properties.) But covariance is of course a statistical notion, not a causal one. If a given collective trait covaries with a collective's fitness, this *may* be because of a causal influence of that trait on fitness, or it may be for some other reason. Another possibility is that the trait-fitness covariance at the collective level is a side-effect, or byproduct, of natural selection acting at the lower level. If so, then there is a 'cross-level product' running from the particle to the collective level [2].

Figure 1: Particles nested within a collective

This conceptualization helps us to isolate the core of levels of selection problem. The key question becomes: *when is a given trait-fitness covariance indicative of selection at the level in question, and when is it a byproduct of selection at some other hierarchical level?* In previous work, I argue that this is the question actually at stake in many debates over the levels of selection, though they are rarely formulated in precisely this way [2]. In essence, the

6

levels of selection problem is about how to determine the hierarchical level(s) at which there is a causal, rather than merely a statistical, link between traits and fitness.

To illustrate the idea of a cross-level byproduct, recall G.C. Williams' example of a herd of fleet deer. Let us elaborate on the example somewhat. Suppose there are two sorts of deer, fast and slow, and that the former are on average fitter than the latter, as they can escape predators more easily. The deer live in herds of size n. Suppose that the fitness of any individual deer depends only on its own running speed, and not on which group it lives in. Suppose that fast and slow deer are distributed among herds at random, so by chance, so the proportion of fast deer (for example), varies from herd to herd. In this situation, there will clearly be a positive covariance between a herd's fitness (which can be defined as the total fitness of the deer within it), and the proportion of fast deer in the herd. Herds in which this proportion is high will be fitter than ones in which it is low. However, this covariance does not reflect a causal relationships of trait on fitness at the herd level; rather, it is a side effect of the fact that at the individual level, there is a causal link between running fast and being fit. So the causal action of natural selection is taking place at the lower level, producing effects that 'filter up' the biological hierarchy, leading to the *appearance* of a causal process of selection at the higher level.[1]

How exactly should the distinction between 'genuine' natural selection at a level and cross-level byproducts be drawn? Some authors have suggested that emergent properties can help distinguish the two. Where the collective trait is aggregate rather than emergent, then any covariance between that trait and fitness can only be a side effect of lower-level selection, on this view; while if the collective trait is emergent, then it *is* capable of causally influencing fitness. This has been called the 'emergent property requirement' on genuine higher-level selection [2]. The requirement has a certain plausibility, and certainly tallies with our intuitions in some cases, but it faces two problems. Firstly, as noted above, the aggregate/emergent distinction, though intuitive, is hard to characterize precisely. Secondly, it represents a substantial metaphysical thesis whose truth one would like some explanation of. However exactly we distinguish emergent from non-emergent properties of collectives,

---

[1] This has the important consequence, often not noticed, that Price's equation is a potentially misleading guide to the levels of selection, if 'selection' is understood causally. This point has been made in the literature by a number of authors [2, 14, 15] but is not widely appreciated.

why should it be that only the former are capable of causally influencing the fitness of a collective? Proponents of the emergent property requirement have not answered this question.

Another idea that has surfaced in this debate, though rarely made fully explicit, is that genuine collective-level selection, which cannot be reduced to selection at lower levels, is in fact impossible. (Both Vrba [12] and Eldredge [13] flirt with an argument which, if taken to its logical conclusion, would have this consequence.) On this view, *any* trait-fitness covariance at the collective level can ultimately be explained from below; so there *cannot* be a causal link between trait and fitness at the collective level as opposed to a cross-level byproduct. One possible argument for this reductionistic conclusion is as follows.

In general, properties of collectives are likely to depend systematically on properties of their constituent particles – this is the principle of 'part-whole supervenience'. So any collective trait Z will be 'realized' by some complex of underlying particle traits; the same is true of collective fitness Y. Therefore, there cannot be a direct causal influence of Z on Y. Any apparent causal link between Z and Y is actually a side effect of causal connections between the respective particle-level traits that realize Z and Y. So genuine collective-level selection, that is irreducible to causal processes acting at the particle level, is impossible; if flies in the face of the determination of collective properties by particle properties. This argument is depicted graphically in Figure 2. The solid arrows and dotted lines represent causation and correlation respectively; the thick vertical lines represent the relation of determination, or supervenience. Let us call this the 'supervenience argument' against the possibility of genuine higher-level selection.

Figure 2: The supervenience argument against higher-level selection

What should we make of the supervenience argument? Clearly it threatens to make cross-level byproducts ubiquitous, for it challenges the very idea of higher-level causation in a hierarchical system. (The argument is analogous to Jaegwon Kim's well-known argument against non-reductive physicalism in philosophy of mind [5].) However, note that the

supervenience argument, if correct, shows only that a trait-fitness covariance at the higher level must be a by-product of *some lower-level causal processes or other*, but not necessarily lower-level *selection*. For the underlying particle characters on which Y supervenes will not necessarily be particle *fitnesses*; they may be characters of any sort. So it does not follow from the supervenience argument that the trait-fitness covariance at the collective level is reducible to particle-level selection.

However, reducibility to lower-level selection is what matters, for evolutionary biologists. For the crucial question is: should we invoke the notion of Darwinian adaptation at the collective level, or only at the particle level? Should we recognise collectives as adapted units, with properties fashioned by natural selection, or not? In the context of the levels of selection debate, this is the issue that matters. So the mere fact, if it is one, that there will always be some lower-level ('micro-causal') explanation of a given higher-level trait-fitness covariance is not to the point; what we are interested in is whether or not there is a lower-level *selective* explanation. For this latter question is what determines the legitimacy, or otherwise, of treating the collectives as adapted units in their own right, rather than as groups of adapted individuals. And this is precisely the bone of contention in the levels of selection controversy.

This does not show that the supervenience argument is incorrect, but only that it is not quite to the point, given the question that evolutionary biologists are interested in. This also helps us diagnose the mistake made by proponents of the emergent property requirement on higher-level selection. The appeal to emergent properties makes some sense as a way of trying to resuscitate 'genuine' higher-level selection from the clutches of the supervenience argument. Indeed this is a standard role played by the emergent property notion in other areas – helping to explain (supposedly) how causation at the higher level can co-exist with part-whole supervenience. It is highly debatable whether emergent properties can succeed in this role, but in the biological context it does not matter. For to repeat, the question we are interested in is not whether some particle-level causal processes or other bear the causal responsibility, but whether particle-level *selection* bears the causal responsibility. It seems probable that these two questions have been conflated by defenders of the emergent property requirement on higher-level selection.

9

Conclusion

What does this leave us? The upshot, I think, is that the debate over causation, hierarchy and emergence, in this particular area of evolutionary biology, raises issues that are fairly *sui generis* to biology, despite the appearance of a common structure with issues discussed in philosophy of mind, metaphysics, and other branches of science. One interesting question is whether this moral generalizes. Are the issues surrounding emergence and causation, in, say, neuroscience, importantly different to the similar-sounding issues that arise in complex systems theory or in statistical mechanics, for example? There is a perennial temptation in philosophy of science to see a common structure in debates with widely different subject matters, and thus to seek an abstract characterization of the issues, applicable across the board. This can be illuminating, but it can also blind us to the subtleties and idiosyncrasies of particular cases. The problem in evolutionary biology I have discussed *can* be characterized in way that makes salient the analogy with the causal exclusion / emergence / reductionism debate in metaphysics and philosophy of mind, but doing so threatens to obscure the biological issues that are at stake.

References

[1] Bedau, M. and Humphreys, P. 2008 *Emergence*, Cambridge MA: MIT Press.

[2] Okasha, S. 2006 *Evolution and the levels of selection*, Oxford: Oxford University Press.

[3] Michod, R. E. 1999 *Darwinian dynamics: evolutionary transitions in fitness and individuality*, Princeton: Princeton University Press.

[4] Maynard Smith, J. and Szathmáry, E. 1995 *The major transitions in evolution*, Oxford: Oxford University Press.

[5] Kim, J. 1998 *Mind in a physical world*, Oxford, Oxford University Press.

[6] Sklar, L. 2000 *Theory and truth*, Oxford: Oxford University Press.

[7] Queller, D. C. 2000 Relatedness and the fraternal major transitions. *Philosophical Transactions of the Royal Society* B 355, 1647-55.

[8] Buss, L. W. 1987 *The evolution of individuality*, Princeton: Princeton University Press.

[9] Sterelny, K. and Calcott, B. 2011 *Major transitions in evolution revisited*, Cambridge MA: MIT Press.

[10] Price, G. R. [1972] Extension of covariance selection mathematics. *Annals of Human Genetics* 35, 485-90.

[11] Williams, G. C. 1966 *Adaptation and natural selection*, Princeton: Princeton University Press.

[12] Vrba, E. 1989 Levels of selection and sorting with special reference to the species level. *Oxford Surveys in Evolutionary Biology* 6, 111-68.

[13] Eldredge, N. 1989 *Macroevolutionary dynamics: species, niches and adaptive peaks*, New York: McGraw-Hill.

[14] Heisler, I. L. and Damuth, J. 1987 A method for analyzing selection in hierarchically structured populations. *American Naturalist* 130, 582-602.

[15] Nunney, L. 1985 Group selection, altruism and structured-deme models. *American Naturalist* 126, 212-230.

Figure Captions

Figure 1: Particles nested within a collective

Figure 2: The supervenience argument against higher-level selection

Short Title

Emergence in Biology