Précis of *Evolution and the Levels of Selection*

SAMIR OKASHA
University of Bristol

The 'levels of selection' question is one of the most fundamental in evolutionary biology, for it arises directly from the logic of Darwinism. As is well-known, the principle of natural selection is entirely *abstract*; it says that any entities satisfying certain conditions will evolve by natural selection, whatever those entities are. (These conditions are: variability, associated fitness differences, and heritability (cf. Lewontin 1970).) This fact, when combined with the fact that the biological world is hierarchically structured, i.e. smaller biological units are nested within larger ones, gives rise to the levels of selection question. For in principle, entities at many different hierarchical levels, above and below that of the 'individual organism', (e.g. gene, chromosome, cell, kin group, colony, lineage, species) can satisfy the requirements for Darwinian evolution. This possibility has long been recognised by biologists, from Darwin himself to contemporary proponents of 'multi-level selection'; and there exist numerous biological phenomena which suggest that it has actually occurred.

Biological altruism, in which one organism performs a behaviour which reduces its own chance of survival/reproduction but benefits that of others, is an example of a phenomenon that, *prima facie*, is indicative of selection occurring at a level other than that of the individual organism. For altruism, by definition, is individually disadvantageous, and yet is common in the animal kingdom. Darwin himself suggested that altruism may have evolved by group-level selection, i.e. groups containing many altruists out-performed groups containing fewer, offsetting the individual cost of behaving altruistically. This suggestion, though controversial, is still taken seriously by many contemporary biologists; and the more general link between altruism and levels of selection remains as intimate as in Darwin's day.

From what has been said so far, the levels question may seem purely empirical. Given that selection *can* operate at many different hierarchical levels, surely it is just a matter of finding out which levels it *does* (or did) operate at? On the face of it, this might seem a straightforward scientific question, resolvable with sufficient empirical data. In fact, however, the literature on levels of selection is up to its neck in conceptual and philosophical issues, which have caused much confusion and continue to do so today. Unless we can agree on what it means for there to be selection at a given hierarchical level, or what the criteria for individuating 'levels' are, on whether selection at one level can ever be 'reduced' to selection at another, on how multi-level selection should be modelled, and on whether there is always 'one true fact' about the level(s) at which selection is acting, then there is little prospect of empirical resolution, however much data we collect.

To get a handle on these conceptual issues, it helps to operate with an abstract mathematical description of evolution by natural selection, applicable to entities of any sort, at any hierarchical level. *Price's equation*, named after the American geneticist George Price, is ideal for this purpose. The equation provides a simple, general way of describing an evolving population, and subsumes all more specific evolutionary models as special cases. Though really just a mathematical tautology, the equation is conceptually invaluable, and reveals something deep about the Darwinian process, namely that *character-fitness covariance* is the essence of natural selection, and that the evolutionary change produced by selection, in any population, is determined by the magnitude of this covariance. This is highly intuitive: if taller organisms leave more offspring than shorter ones, i.e. if the character 'height' covaries positively with fitness, then we expect the average height of the population to increase, i.e. directional evolutionary change to occur. There turns out to be an intimate link between Price's equation and Lewontin's tripartite analysis of the conditions required for Darwinian evolution, mentioned above.

Price's equation has a special significance for the levels of selection that goes beyond its great generality, for it lends itself naturally to a description of *multi-level* selection, as Price himself realised (cf. Price 1972). Suppose we have a hierarchically structured population, consisting of 'particles' nested within 'collectives' as in Figure 1 below. (For example, the particles could be social organisms and the collectives the colonies they form; the particles could be cells and the collectives the multi-celled organisms they make up; or the particles could be genes and the collectives the genomes they reside in.) In the face of such hierarchical structure, Price's basic equation can be expanded into a 'multi-level' format, which allows the combined effects of the two levels
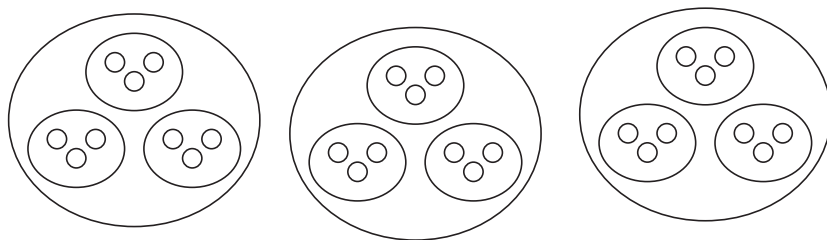
Figure 1: Hierarchical Organization

of selection (particle and collective) on the overall evolutionary change to be represented in a single schema. This expansion of Price's equation played a key role in the genesis of modern multi-level selection theory, and informs the way many contemporary biologists think about the levels of selection (cf. Hamilton 1975).

Graphical depictions of hierarchical organization in biology, as in Figure 1 above, are quite common; however it is not always clear what actual biological relation(s) correspond to the abstract relation of containment depicted in the figure. There are two main possibilities: firstly that *interaction among the particles* is what binds them into a collective, and secondly that the particles within a given collective are genealogically related. These two relations—ecological interaction and shared ancestry—give rise to different types of hierarchical structure. The 'interactionist' conception of part–whole structure fits many aspects of the levels of selection question, but not all. Selection can also operate on entities, such as whole species, that are genealogically-defined.

What exactly does multi-level selection involve? At a single level, evolution by natural selection requires character (or trait) differences, associated differences in fitness, and heritability; so selection at multiple levels presumably requires these features to be present at more than one hierarchical level. For example, if group-level selection is to occur, in addition to individual selection, then presumably groups as well as individuals must exhibit character differences, associated fitness differences, and heritability. This raises an overarching question: what is the relation between the characters, fitnesses, and heritabilities at the different levels? For example, does a collective's character depend somehow on the characters of its constituent particles? Is the fitness of a collective simply the aggregate of the particle fitnesses? These metaphysical-sounding questions are rarely discussed explicitly by biologists, but much theorising about levels of selection implicitly presumes certain answers to them.

Two key distinctions help clarify how natural selection works in a hierarchical world. The first is a distinction between two types of multi-level selection, introduced by Damuth and Heisler (1988). In multi-level

selection type 1 (MLS1), the particles are the 'focal' units, i.e. the units whose demography gets tracked; the collectives in effect constitute part of the particles' environment. The point of an MLS1 analysis is to explain the changing frequency of particle-types, in the overall meta-population of particles. In multi-level selection 2 (MLS2), the particles and collectives are both focal units; in an MLS2 analysis, the frequencies of both the particle-types and the collective-types are tracked. Following Damuth and Heisler, I argue that MLS1 and MLS2 are both bona fide types of multi-level selection. There are important logical differences between them, in particular in how the notion of 'collective fitness' is defined in each; both can be usefully analysed using the Price equation.

The second key distinction is between what I call 'direct' or 'real' selection at a given level and a mere 'cross-level byproduct'. Direct selection at a level occurs when there is a character-fitness covariance at that level which arises because there is a causal link between the character and fitness, e.g. the character enhances fitness. A cross-level byproduct arises when a character-fitness covariance at a given level does not arise from a causal link at that level, but is a side-effect, or byproduct, of direct selection at a *different* hierarchical level. For example, suppose there are two types of individual, A and B, living in a group-structured population. The fitness of an individual depends solely on its own type; suppose that As are intrinsically fitter than Bs. Suppose further that the proportion of A types in a group varies across the groups. Then, there will be a positive covariance between the group character 'proportion of A types' and group fitness (understood as total individual fitness.) However, there is no direct selection at the group level, in this example; rather, individual-level selection is doing all the causal work, and is producing effects that 'filter up' the biological hierarchy, creating the mere *appearance* of a group-level selection process.

The distinction between direct selection and cross-level byproduct brings to the fore the causal dimension of Darwinian theory. It is obvious that Darwinian explanations are causal; to attribute a trait's spread in a population to natural selection is to advance a hypothesis about what caused it to spread. It is equally obvious that a trait can correlate with fitness without causally affecting fitness, e.g. if the trait is correlated with a second trait that *does* causally affect fitness. Biologists sometimes capture this point by distinguishing between 'direct' and 'indirect' selection on a trait; in Sober's well-known terminology, it is the distinction between 'selection of' and 'selection for' (Sober 1984).

In essence, the direct selection / cross-level byproduct distinction arises from extrapolating the selection of/for distinction, or the

direct/indirect selection distinction, to a hierarchical setting. The key point is that direct selection at any one level may have effects at other hierarchical levels, higher and lower, *and one of these effects can be a character-fitness covariance at that other level*. This point is related to (though not the same as) G.C. Williams's well-known point that 'group adaptation' is different from 'fortuitous group benefit' (Williams 1966), a point that many evolutionists are sympathetic to. However, the importance of the point has not always been appreciated in the literature, particularly by fans of the 'Price equation' approach to multi-level selection, and has not previously been given a systematic formal and philosophical analysis.

The notion of a cross-level byproduct suggests a way of setting up the levels of selection problem that is particularly sharp. The key question becomes: *when is a character-fitness covariance indicative of direct selection at the level in question, and when is it a by-product of direct selection acting at a different level?* I suggest that this is the question actually at stake in much of the literature, and helps us make sense of various proposals for how to determine the 'real' level(s) of selection in a given circumstance. For example, Hull's (1981) 'interact as a cohesive whole' criterion, Wimsatt's (1980) and Lloyd's (1988) 'additivity criterion', Vrba's (1989) 'emergent character' criterion, Brandon's (1990) 'screening off' criterion, and Gould's (2002) 'emergent fitness' criterion, can all be understood as attempts to specify when a given character-fitness covariance reflects direct selection at the level in question, and when it is a cross-level byproduct. This enables these proposals to be placed in a common framework, and their validity assessed.

One major source of philosophical interest in the levels of selection debate is the opposition between realism and pluralism (or convetionalism). In this context, pluralism is the idea that in some or all cases, there is no objective fact about the level(s) at which selection is acting; different answers to the question are equally correct. Realism, by contrast, says that there *is* always an objective fact about the level(s) of selection. Though pluralism is popular among philosophers, realism is arguably the natural default position and is assumed by most biologists. Three different arguments for pluralism about levels of selection can be discerned in the literature. The first derives from a non-realist account of causation; the second from the indeterminacy of hierarchical organization; and the third from the existence of mathematically interchangeable descriptions. I argue that a philosophically interesting form of pluralism is defensible only in very specific circumstances.

A second source of philosophical concern lies in the opposition between reductionist and non-reductionist approaches to the levels problem. However, at least three different concepts of reductionism

have featured in the levels of selection debate, not always sharply distinguished from each other. The first is the general idea that properties of wholes should be explained in terms of properties of their parts; this is related to the doctrine of methodological individualism in social science. The second is the idea, associated with G.C. Williams in particular, that lower levels of selection are explanatorily preferable to higher levels. The third is the idea that selection at one hierarchical level may be "reducible" to selection at a different level, e.g. group selection may be reducible to individual selection. This idea is related to the notion of a cross-level byproduct, discussed above. The three concepts of reductionism are logically independent of each other.

To this point, my treatment of the levels of selection issue has been deliberately abstract, framed in terms of 'particles' and 'collectives', with no reference to any particular level of biological organization. However, many theorists hold that the *gene* level is in some ways privileged; this position has been articulated most clearly by Richard Dawkins (1976), who argues that all evolutionary phenomena can be understood in terms of competition between 'selfish genes' for increased representation in the gene pool. How exactly this 'gene's eye view' of evolution is related to the levels of selection problem, and in particular how it relates to multi-level selection theory, is not a simple matter.

This is partly because the genic approach suffers from a certain ambiguity in status: sometimes it is presented as an empirical thesis about the course of evolution, at other times as a heuristic perspective for thinking about evolution. This ambiguity can be resolved by distinguishing *genic selection*, which is a causal process, from the *gene's eye viewpoint*, which is a perspective. Genic selection occurs when there is selection between the genes within a single organism, or genome; it is thus a distinct level of selection of its own. By contrast, a gene's eye view can be adopted on selection processes occurring at various hierarchical levels, not just the genic level. This distinction between the process of genic selection and the gene's eye viewpoint enables much of the controversy surrounding 'selfish gene' thinking to be resolved.

Understood this way, genic selection is relatively rare—it occurs only in cases of intra-genomic conflict, where the genes within a single genome have divergent evolutionary interests; such genes are sometimes called 'outlaws' or 'ultra-selfish genes'. Genic selection can lead a gene to increase in frequency even if the gene has harmful effects on its host organism, and so is counter-selected at the organismic level. Therefore, the dynamics of outlaw genes typically involve selection at multiple hierarchical levels. This is somewhat ironic, given the widespread (though mistaken) tendency to see multi-level or hierarchical views of selection as somehow opposed to a genic view.

Historically, the question of *group selection*—where 'group' means a group of multi-celled organisms—has been one of the mainstays of the levels of selection debate. As is well-known, the concept of group selection has enjoyed a chequered career in evolutionary biology. It was widely dismissed in the 1960s and 1970s, before making something of a comeback in the 1990s, when theorists such as Sober and Wilson (1998) argued that the traditional anti-group selection consensus was flawed. The controversy unabated continues to this day (cf West *et al.* 2008, Gardner and Grafen 2009). In part, the ongoing controversy reflects disagreement about empirical matters, but it also has a conceptual dimension. There is widespread confusion about the relation between kin and group selection, the 'correct' way of modelling group selection, and the relation between group selection and altruism.

Many of these issues can be resolved, or at least clarified, using the abstract framework for thinking about levels of selection outlined above. At root, the controversy can be traced to the fact that two different requirements have been taken as necessary for a process to count as group-level selection (cf. Okasha 2004). The first is variance in group fitness—some groups must have higher reproductive output than others. The second is a 'group effect' on individual fitness—the fitness of an individual must be affected by its interaction with other group members. (Where this second condition is not satisfied, any variance in group fitness is a side-effect, or byproduct, of selection at a lower level.) The two main ways of modelling group selection in the literature each respect one of these requirements but violate the other.

A separate facet of the levels of selection debate is the controversy over species selection. This is the idea, defended by biologists such as S.J. Gould (2002), that a selection process might operate on whole species, preserving the ones best able to survive and ∕ or reproduce (i.e. produce daughter species). This might seem like a higher-level analogue of group selection, where the 'group' is the whole species, but in fact there are important logical differences between group and species selection. This is because most models of group selection are of the MLS1 variety, i.e. the aim is to explain the evolution of an individual trait, often altruism, in a group-structured population, while models of species selection are of the MLS2 variety, the aim being to explain the changing frequency of different types of species, not individuals.

In recent biology there has been considerable interest in what have come to be called 'major evolutionary transitions' (cf. Maynard Smith and Szathmáry 1995, Michod 1999). These transitions occur when a number of free-living biological units, capable of surviving and reproducing alone, become integrated into a cooperative whole, generating a new level of biological organization. (Think for example of the

evolution of modern multi-celled organisms from single-celled ances-
tors, or of social insect colonies from solitary insects). Such transitions
have occurred numerous times in the history of life on earth, giving rise
to the modern biological hierarchy. Clearly, evolutionary transitions
create the potential for conflict between levels of selection, for selection
between the smaller units may disrupt the well-being of the collective.

I believe that the traditional levels of selection question has been
subtly transformed by recent work on evolutionary transitions. In tra-
ditional discussions, such as Lewontin's well-known 1970 treatment,
the existence of the biological hierarchy was taken for granted; the
question was about selection and adaptation at pre-existing hierarchical
levels. But the evolutionary transitions literature is concerned with the
origins of hierarchical organization itself; this requires a "diachronic"
rather than a "synchronic" formulation of the levels of selection ques-
tion. This new diachronic perspective requires that hierarchical organi-
zation be "endogenized" by evolutionary theory, and itself be given a
Darwinian explanation.

The shift from a synchronic to a diachronic formulation of the levels
question forces a re-think of many theoretical issues, including the rela-
tion between the two types of multi-level selection, the significance of
'emergent properties', the role of kinship, and the notion of a biological
individual. Interestingly, it turns out that many of the 'problem cases' in
the earlier levels of selection literature, e.g. cases where theorists disagree
about how the level(s) of selection should be identified, or about whether
a single or multiple levels of selection are at work, turn out to represent
intermediate stages in evolutionary transitions. This makes the disagree-
ments surrounding their status more intelligible, for borderline cases of
part–whole structure are inevitable in an evolutionary transition, given
the gradualness of the Darwinian process. That such disagreements can
persist, even when the basic empirical facts are not in dispute, is readily
understandable once we adopt a diachronic approach.

More generally, the conceptual issues that form the core of the tradi-
tional levels of selection debate, in both the biological and philosophi-
cal literatures, are subtly transformed when we move from a
synchronic to a diachronic formulation of the levels question, as we
must do if we are to understand the evolutionary transitions in Dar-
winian terms. But the transformation is not so drastic that the tradi-
tional discussions lose all their relevance.

References

Brandon, RN 1990. *Adaptation and Environment*. Princeton: Princeton
University Press.

Damuth, J and Heisler, IL 1988. 'Alternative Formulations of Multi-level Selection'. *Biology and Philosophy* 3: 407–30.

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

Gardner, A and Grafen, A 2009. 'Capturing the Superorganism: towards a formal theory of group adaptation'. *Journal of Evolutionary Biology* 22(no. 4): 659–71.

Gould, SJ 2002. *The Structure of Evolutionary Theory*. Cambridge MA: Harvard University Press.

Hamilton, WD 1975. 'Innate Social Aptitudes in Man: an Approach from Evolutionary Genetics'. In: R Fox (ed.) *Biosocial Anthropology*. New York: Wiley, 133–55, reprinted in his *Narrow Roads of Gene Land*, volume 1, New York: W. H. Freeman, 315–52.

Hull, DL 1981. 'Units of Evolution: a Metaphysical Essay'. in UJ Jensen and R Harre (eds.) *The Philosophy of Evolution*. Brighton: Harvester Press 23–44.

Lewontin, RC 1970. 'The Units of Selection'. *Annual Review of Ecology and Systematics* 1: 1–18.

Lloyd, EA 1988. *The Structure and Confirmation of Evolutionary Theory*. New York: Greenwood Press.

Maynard Smith, J and Szathmáry, E 1995. *The Major Transitions in Evolution*. Oxford: Oxford University Press.

Michod, RE 1999. *Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality*. Princeton: Princeton University Press.

Okasha, S 2004. 'Multi-level Selection and the Partitioning of Covariance: a Comparison of Three Approaches'. *Evolution* 58(no. 3): 486–494.

Price, GR 1972. 'Extension of Covariance Selection Mathematics'. *Annals of Human Genetics* 35: 485–90.

Sober, E 1984. *The Nature of Selection*. Chicago: Chicago University Press.

Sober, E and Wilson, DS 1998. *Unto Others: the Evolution and Psychology of Unselfish Behavior*. Cambridge MA: Harvard University Press.

Vrba, E 1989. 'Levels of Selection and Sorting with Special Reference to the Species Level'. *Oxford Surveys in Evolutionary Biology* 6: 111–68.

West, SA, Griffin, AS and Gardner, A 2008. 'Social Semantics: how useful has group selection been?' *Journal of Evolutionary Biology* 21: 374–385.

Williams, GC 1966. *Adaptation and Natural Selection*. Princeton: Princeton University Press.

Wimsatt, W 1980. 'Reductionist Research Strategies and their Biases in the Units of Selection Controversy'. in T Nickles (ed.) *Scientific Discovery: Case Studies*. Dordrecht: Reidel, 213–59.