

- Israel, D. and J. Perry, 1990, "What Is Information?", in P. Hanson (ed.), 1990, *Information, Language and Cognition*, University of British Columbia Press, Vancouver, pp. 1–19.
- O'Rourke, M. and C. Washington (eds.), 2007, *Situating Semantics: Essays on the Philosophy of John Perry*, MIT Press, Cambridge, Mass.
- Perry, J., 2006, *Referencialismo crítico: La teoría reflexivo-referencial del significado*, trans. K. Korta and R. Aguerri, CSLI Publications, Stanford. —
- , 2002, *Identity, Personal Identity, and the Self*, Hackett Publishing, Indianapolis.
- , 2001, *Knowledge, Possibility and Consciousness*, MIT Press, Cambridge, Mass.
- Taylor, K., 1995, "Meaning, Reference, and Cognitive Significance", *Mind and Language*, vol. 10, pp. 129–180.
- Wettstein, H., 1991, *Has Semantics Rested on a Mistake?*, Stanford University Press, Stanford.
- , 1986, "Has Semantics Rested on a Mistake?", *The Journal of Philosophy*, vol. 83, no. 4, pp. 185–209 (reprinted in Wettstein 1991).

LENNY CLAPP

Instituto de Investigaciones Filosóficas
 Universidad Nacional Autónoma de México
 lclapp@filosoficas.unam.mx

Samir Okasha, *Evolution and the Levels of Selection*, Clarendon Press, Oxford, 2006, 263 pp. + xii.

Okasha's book on the levels of selection is the most comprehensive philosophical discussion of the subject in relation to the idea of major evolutionary transitions. The following is a critical review of his discussion of some problems and concepts involved in understanding multilevel selection as a causal process. The book begins with an argument in favor of multilevel selection based on the existence of biological units at various hierarchically organized levels, with lower level particles nested within collectives (genes in cells, cells in organisms, organisms in groups). Perhaps a biological world with a hierarchy of levels of biological units is possible without selection operating simultaneously at the different levels. Our world, however, seems to have undergone processes that began with free ranging particles in interaction and ended up in collectives with internal cohesion, functioning as adaptive units. These processes involved selection and involved the leap from genes to genomes, from cells to multi-cellular

organisms and from organisms to social groups (pp. 16–17). The transition to new hierarchical levels implied conflicts between adjacent levels, examples of which persist today: some traits seem not to favor the multi-cellular organism, but give advantage to lower level units (e.g., meiotic drive) or to higher ones (classic group selection overriding individual selection). In our world the hierarchy of levels implies multilevel selection.

Collective Fitness₁

Multilevel selection (MLS) requires, “character differences, associated differences in fitness and heritability” at more than one hierarchical level (p. 47). An important goal of the book is to give a precise account of character, fitness and heritability at each level and of their relation. As a preliminary step, Okasha discusses the sort of collective that can count as a level in the hierarchy. He argues that fitness-affecting interactions between free ranging particles, capable of surviving and reproducing on their own, is enough for a collective to count as such and rejects more demanding criteria: functional organization, or conflict-reducing mechanisms. These phenomena evolve in virtue of selection. Acting on loose collectives that are like ancestors of collectives with full individuality, MLS explains their evolution and cannot presuppose them (pp. 41–46).

Okasha emphasizes the importance of the distinction between two senses of MLS —MLS₁ and MLS₂— noted also by other authors (p. 56). The distinction is related to two ways in which the fitness of collectives can be understood: collective fitness₁ and collective fitness₂. Collective fitness₁ defines the fitness of the collective as the average value of the fitness of its particles. Collective fitness₂ defines it as the number of offspring collectives. I think it is useful to introduce a parallel distinction between two types of collectives: collectives₁ and collectives₂, accepting grey zones between them. This distinction is implicit in his discussion. Parent-offspring relationships, and presumably also functional organization and repression of internal competition exist clearly in collectives₂, but not quite in collectives₁ (p. 58). Multi-cellular organisms, with functional internal organization and sexual reproduction, for example, are clear instances of collectives₂. Collectives₁ look rather like aggregates and attract attention to the properties of the particles, which are the focus of MLS₁; and notably, to the social properties of these particles, for their evolution “could be the first step in explaining the existence of cohesive collectives, whose constituent particles work for the good

of the whole” (p. 59). Collectives₁ are also evolutionary ancestors of collectives₂. This is implicit also in Michod’s theory on the evolutionary transition to multi-cellular organisms.

Okasha’s explanation of collective fitness₁ deserves some comment. “Clearly collective fitness₁ is aggregate —it arises *directly* from the character of the particles within the collective. Fix the fitness of every particle [...] and the fitness₁ of every collective is fixed too” (pp. 55–56, my emphasis). There is an emergent aspect in collective fitness₁, which Okasha underplays in passages like this one. When collective fitness₁ is defined only in terms of average particle fitness, two different scenarios are being mixed under the same category. In one scenario, particle fitness is independent of group membership. If *z* is the fitter particle character, then collectives with a higher proportion of *z* are fitter, but the fitness of particle character *z* does not change across collectives. Alternatively, collectives with a higher proportion of *z* are fitter, but the fitness of *z* also changes across collectives, the fitness of *z* being higher in collectives with higher proportions of *z*. Collective fitness depends in the second case on something else besides particle fitness, because particle fitness depends on collective character in turn. Of course, Okasha is aware of this dialectic. But as I see it, it implies that collective fitness₁ cannot be defined *solely* in terms of average particle fitness.

The previous reflection questions his claim about the reductionist character of MLS₁ and of contextual analysis: contextual analysis “is intrinsically reductionist, simply because it deals with MLS₁”. It aims at understanding complex group level phenomena “in terms of behavior of individuals” (p. 93). “MLS₁ necessarily involves the bottom-up mode of explanation [...] light is shed into a collective-level feature by explaining the evolution of the underlying particle characters on which the feature depends” (p. 139). This is going too fast, for one could as well say that MLS₁ explains the properties of particles top-down, from properties and processes at the collective level. The same caution applies to claims that collective character and collective fitness “are defined in terms of”, or “are logical constructs of” (p. 91) particle character and fitness. Other components of the definition, like dependence of particle fitness on group character, hinder a reduction of collective fitness₁ to average particle fitness.

By-Products Up and Down the Hierarchy

Many attacks against MLS have exploited the idea that collective selection is merely apparent and can be reduced to particle selec-

tion. This means that causality operates only at the particle level. In contrast, collective selection takes place when a higher-level causality impacts on lower level processes. In Okasha's terminology, there are two types of by-products. One of them occurs when an observed covariance between collective character and collective fitness is a side-effect of character/fitness covariance at the particle level. In contrast, when the covariance at the collective level is independent, there is a by-product effect going from collective fitness to particle fitness (group effects). In this case, at least part of the character/fitness covariance observed at the particle level is a by-product of covariance at the collective level. The existence of by-products going down the hierarchy, i.e., the influence of group character on particle fitness, is the hallmark of MLS processes.

In order to test collective selection on a given system, one has to measure and calculate with the appropriate equations. Okasha initially explores the Price equation as the mathematical expression of natural selection and interprets it in terms of causal decomposition. He then uses it for representing selection at two levels with strict nesting (a given particle cannot be part of two collectives at the same time). It represents a way of partitioning the change due to selection acting at different levels. But the possibility of spurious correlations and by-products, where selection happening at lower levels filters up and produces an illusion of selection happening at higher levels, leads him to acknowledge the superiority of another method of partitioning the total evolutionary change: the contextual approach. The contextual approach acknowledges that there is always a by-product effect of the particle level at the collective level (p. 98). This approach isolates the by-product and makes it possible to judge whether a real collective effect is present. This argument for the contextual approach is one of the peaks of Okasha's book.

With these concepts Okasha then examines several proposed criteria to establish whether selection at the collective level is really happening. His partial condemnation of the additivity criterion when applied to collectives above the genic level deserves a critical comment. The additivity criterion establishes a distinction between linear and non-linear dependence of *collective fitness* on the number or proportion of particles that benefit the collective. If the dependence is linear, the additivity criterion judges that there is no selection at the collective level. Okasha notes that it is common for MLS_1 models to represent group fitness as a linear function of number or proportion of altruistic particles. But he also notes, importantly, that those models are relying on the Price approach, which is blind to cases where

collective fitness is a by-product of particle fitness (pp. 85ff). It is therefore possible to think that the additivity criterion is designed to exclude these MLS_1 models as representing cases where collective fitness is a mere by-product. Okasha then examines whether the additivity criterion may be legitimized when judging its relevance from the perspective of the contextual approach. His result is a partial vindication, and partial condemnation, of additivity, where I think he should vindicate it totally. He is misled by a subtle misinterpretation of additivity. He reformulates it as containing two claims:

- 1) If there is collective-level selection, there must be non-additive variance in collective fitness.
- 2) If there is no collective-level selection, then any variance must be additive (pp. 116–17).

He rejects 1) because, in soft selection, there is no variance in collective fitness and nonetheless the contextual approach detects selection at the collective level. But the additivity criterion is better represented by 1’):

- 1’) If there is collective-level selection, then any variance in collective fitness must be non-additive.

Comparing 1’) to 1), it seems that his partial condemnation of additivity comes from misplacing the modal “must” in his reformulation. Additivity does not mean to say that for selection to operate at the collective level there *must* be variance in collective fitness and that it is non-additive. It only says that *if* there is any variance in collective fitness it must be non-additive for selection to operate at that level; or, it can be easily reformulated to say this, and to accommodate for soft selection with a clause to the effect that the fraction of collective fitness claimed by types that benefit the group increases as a nonlinear function of the proportion of those types.

It is possible to explain where MLS_1 models go wrong when they treat collective fitness as a linear function of the number of altruists. The graphs that thus depict collective fitness in Kerr and Godfrey-Smith (2002, pp. 488, 491),¹ for example, are synchronic snapshots of what happens across collectives. If you want to depict MLS_1 , you

¹ Kerr, B., and P. Godfrey-Smith, 2002, “Individualist and Multi-Level Perspectives on Selection on Structured Populations”, *Biology and Philosophy*, vol. 17, pp. 477–517.

further need to represent what happens to the collectives across generations. The basic assumption is that, across generations, collectives re-compose their proportions of altruists. This is equivalent to representing altruists with a bias in the direction of their altruist donations, namely, to benefit mainly altruists. Without re-composition or bias, the differences across collectives are mere by-products of average particle fitness. In order to visualize this in term of equations, consider a simple model where collectives are composed only of two types: a (altruist) and s (selfish). The baseline fitness is 1; assume a confers b units of fitness to every other member of the collective, including itself, and pays a cost of c (if $c > b$, a is strongly altruistic). N is the number of members in a collective and n the number of altruists. In this case both the fitness of particles ($F(a, s)$) and the fitness of collectives ($F(C)$) increase linearly with the number of altruists:

$$F(a) = 1 - c + bn$$

$$F(s) = 1 + bn$$

$$F(C) = (1 - c + bn)n + (1 + bn)(N - n)$$

$$F(C) = n - cn + bn^2 + N - n + bNn - bn^2$$

$$F(C) = bNn - cn + N, \text{ which is a linear function of } n.$$

Assume now that altruists benefit only other altruists (all the parameters remaining equal).

$$F(a) = 1 - c + bn$$

$$F(s) = 1$$

$$F(C) = (1 + bn - c)n + (N - n)$$

$$F(C) = n + bn^2 - cn + N - n$$

$$F(C) = bn^2 - cn + N, \text{ which is a non-linear function of } n.$$

The difference between both calculations is this: in the first, the probability that an altruistic donation goes to a given type is equal to the proportion of this type in the collective. In the second case, in contrast, there is a bias for directing donations at collective-benefiting types. This bias determines that collective fitness is a non-linear function of the number of collective-benefiting types in the population.

Without bias or re-composition, even though collectives with more altruists have higher fitness, all selection is happening at the particle level and altruists will die out. The second case represents the re-composition of collectives as a bias of donations toward altruists. With this assumption, collective fitness immediately jumps from a linear into a non-linear function of the number of altruists in the group.

MLS and Evolutionary Transitions

The influence of group character on particle fitness is the hallmark of both MLS_1 and MLS_2 . In Okasha's terminology, the MLS approach is legitimate when there are by-product effects going from collective fitness₁ or collective fitness₂ to particle fitness (group effects). Okasha, however, stresses that, in MLS_2 , collective fitness₂ and the fitness of the particles are disconnected, at least as a matter of definition. He underlines the difference saying that in MLS_2 "selection at each level leads to a different type of evolutionary change, measured in different units" (p. 57); and also: "in MLS_2 , the evolutionary changes caused by the two levels of selection are incommensurable" (p. 75). Moreover, he presents no MLS_2 equation to add the effects of selection at different levels, but, rather, two different equations, one for each level (p. 74). According to this, MLS_2 is not about two or various levels contributing to one selection process.

However, it seems to me that the theory of major evolutionary transitions postulates a causal connection between collectives₁ and collectives₂, and therefore, between collective fitness₁ and collective fitness₂. Okasha seems to believe that when the transition is complete, the connection is severed. I do not believe this happens. Some phenomena seem to confirm that the difference between MLS_1 and MLS_2 is not sharp. Take for example meiotic drive. Organisms with sexual reproduction are clear instances of collectives₂, so collective fitness₂ applies clearly to them. However collective fitness₂ is not independent of distorter alleles or their suppressors, nor is the fitness of these particles unaffected by what happens to the collectives (organisms). Distorter particles disappear because sick organisms disappear. Since we are counting particles, Okasha says this is MLS_1 (p. 69) and a case where MLS_1 operates on an emergent character (p. 57). But alternatively, phenomena like this one suggest that MLS_1 remains a possible standpoint even when MLS_2 processes are already in place. MLS_2 cannot completely sever the connection between collective fitness₂ and collective fitness₁ or "decouple" the former from the latter. A high

proportion of particles that benefit the collective —or high collective fitness₁— is necessary to sustain high collective fitness₂.

To visualize this fully, it is necessary to avoid a mistaken interpretation of Michod's concept of fitness-decoupling. At some point, Okasha seems to understand that fitness-decoupling happens when an MLS₂ process emerges from an MLS₁ process (p. 232). But in fact, fitness-decoupling, as addressed by Michod and collaborators, refers to the emergence of a genuine collective from what is only the appearance of one. Okasha himself notes this when he says: "Michod and Roze's discussion brings out an interesting link between fitness decoupling and the concept of a cross level by-product" (p. 235). Before fitness-decoupling takes place, collective fitness is a cross-level by-product and does not support an MLS process of any type. Commenting on Michod, Okasha shifts, apparently inadvertently, from the idea of decoupling from a by-product effect to the idea of decoupling from an MLS₁ process, for example on a passage on pages 237–238.

In fact, the theory of evolutionary transitions implies, in my view, that no decoupling can take place between collective fitness₂ and collective fitness₁. This is just a consequence of the fact that MLS₁ explains the origin of collectives₂ starting from collectives₁. A multicellular creature where "differences in organismic fitness [...] arise solely through differences in cell fitness" (pp. 234–235) is not yet a genuine collective₁. Its fitness is a mere by-product of the fitness of its particles. The fitness of a genuine collective₁, even though it does not produce offspring collectives, cannot be fully explained by the fitness of its particles alone.

There is something genuinely "emergent" in collectives₁, as Damuth and Heisler have noted. Okasha also comments on this. The relation between collective character and collective fitness₁ is emergent, because "it cannot be accounted for by the character-fitness relation at a lower hierarchical level" (p. 119). Recall that in MLS₁ collective character is addressed as average particle character, and collective fitness as average particle fitness. Why then can the relation between particle character and particle fitness not explain the relation between collective character and collective fitness? The reason is that there is no independent covariance of particle character to fitness that explains collective fitness. A particle character does not have a fitness assigned to it independently of the collective it is in. The fitness of the particle is affected by some property of the collective, usually the frequency of a character that benefits the

collective. Collective fitness emerges as a function of a property of the collective, not only as a function of its particle characters.

Okasha's book provokes readers to think the issues further, and this is probably the best quality a work can have. One form of repaying the author's effort is to show that his or her work has this quality. I hope I have been able to pay such a tribute to this book.

ALEJANDRO ROSAS

Konrad Lorenz Institute for Evolution and Cognition Research/
Philosophy Department
Universidad Nacional de Colombia
arosasl@unal.edu.co

John Dupré, *Humans and Other Animals*, Oxford University Press, Nueva York, 2002, 272 pp.

Quizá la obra por la cual John Dupré es más conocido es *The Disorder of Things* (1993),¹ donde defiende sagazmente un nominalismo en torno a las clases naturales que críticamente hacía ver la imposibilidad de aplicar en la biología un recuento sobre ellas a la manera de Kripke y Putnam. Estos autores afirmaban que tales clases podían analizarse por medio de la identificación de un *marcador sintáctico*—e g., “ballena”— que venía asociado a un *marcador semántico*—el conjunto de las ballenas— mediante un *estereotipo*—las ballenas son mamíferos piciformes de gran tamaño— que era empleado por los hablantes competentes de una lengua para identificar los objetos que podían ser denotados utilizando tal marcador sintáctico. En un primer momento, esta asociación se establecía por medio de múltiples ostensiones; sin embargo, el desarrollo de la ciencia iba a encontrar un cuarto componente, la *referencia*, que en términos generales consistía en la dilucidación de la estructura causal que justamente explicaba por qué los objetos referidos por dicho marcador sintáctico podían ser identificados con ayuda de tal estereotipo—en nuestro ejemplo, esto sería equivalente al reconocimiento de la monofilia del grupo de los cetáceos así como a la identificación de tendencias evolutivas comunes al grupo—. Esta visión parecía prometer que la ciencia tendería a ir refinando los términos propios de los lenguajes naturales

¹ *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*, Harvard University Press, Cambridge, Mass.