

# The causal homogeneity of biological kinds

Michael Esfeld

University of Lausanne, Department of Philosophy &  
Centre romand for logic, history & philosophy of science

Michael-Andreas.Esfeld@unil.ch

(forthcoming in *History and Philosophy of the Life Sciences* 27 (2005), Dec. issue)

## Abstract

The aim of this paper is to show that biological kinds can be causally homogeneous, although all biological causes are identical with configurations of physical causes. The paper considers two different strategies to establish that result: the first one relies on two different manners of classification (according to function and according to composition); the other one exploits the idea of biological classifications being rather coarse-grained, whereas physical classifications are fine-grained.

## Keywords

biological kinds, causation, functional properties, homogeneity, laws, reduction, supervenience

### 1. *The problem*

The following two claims are well founded:

- 1) Biological kinds are homogeneous kinds: there is a homogeneous pattern of biological effects that characterizes all and only the tokens that come under a given biological type or kind. (I shall use the terms “type”, “kind” and “property” in a synonymous way).
- 2) Biological causes are identical with physical causes. That is to say, any biological token that is part of a causal chain is identical with a configuration of physical tokens. (I shall, without further argument, take property tokens to be the relata of causal relations).

Why is there a problem in holding true the conjunction of these two claims? Recall what is known as the causal exclusion problem as set out by Jaegwon Kim (1998, 38–47; 2005, 36–45). The physical realm is causally complete. Let us call this principle *completeness*: for every physical token  $p$ , insofar as  $p$  has a cause, it has a sufficient physical cause. Otherwise, physics would not be able to explain any physical token  $p$  (insofar as  $p$  has an explanation at all), and the physical laws would not cover any physical token  $p$  (insofar as  $p$  comes under a law at all). Given the state of the art of physics, completeness is a well-established philosophical principle (Papineau 2002, appendix).

Completeness motivates endorsing *supervenience*, given that everything there is in the world developed from a physical basis during cosmic evolution. In its most general form, supervenience is the claim that “any world which is a *minimal* physical duplicate of our world is a duplicate *simpliciter* of our world” (Jackson 1998, 12). That is to say, if applied to the level of fundamental physical tokens: if you imagine the entire distribution of the fundamental physical tokens over the whole of space-time duplicated – in the sense that you copy only that distribution without adding anything that does not exist in our world –, thus creating a *minimal* physical duplicate of our world, the result is that *everything* that there is in our world is duplicated.

This is a thesis of global supervenience. It pins down a way in which everything that there is in the world depends on the fundamental physical level. But it does not establish any specific dependence between particular physical tokens and, say, particular biological tokens. In order to get to specific dependence claims, one has to put forward local supervenience claims – such as the thesis that if you duplicate the physical and the chemical properties of an organism, you thereby create a double of all its properties, including its biological ones. The problem with such supervenience claims is obvious: the biological properties of an organism depend not only on its physio-chemical constitution, but also on its environment. Fitness is an evident case in point. However, it seems that the salient environment is rather limited. Let us therefore take for granted a local supervenience thesis to the effect that all the biological properties of an organism supervene on its physical properties as well as the physical properties of its environment (Rosenberg 1978).

Biological tokens are causally efficacious: every biological token is part of a causal chain, producing some effects. Suppose that a biological token  $b_1$  brings about another biological token  $b_2$ . Given the supervenience of the biological on the physical, if  $b_1$  is to cause  $b_2$ , it has to cause its supervenience base as well. Consequently,  $b_1$  can only cause  $b_2$  by bringing about a configuration  $p_2$  of physical tokens that is a sufficient physical condition for  $b_2$ . However, according to completeness – which constitutes the main reason for endorsing supervenience –,  $p_2$  has a complete physical cause, that is, another configuration  $p_1$  of physical tokens. Assume that  $p_1$  is sufficient as a supervenience base for  $b_1$ . Furthermore, by causing  $p_2$ ,  $p_1$  causes as well all that supervenes on  $p_2$ , that is,  $b_2$ . We thus have two complete causes of  $b_2$ , namely  $b_1$  and  $p_1$ , as well as two complete causes of  $p_2$ , namely  $p_1$  and  $b_1$ . Consequently, given supervenience, the physical causes seem to exclude – or at least to pre-empt – biological causes or any other higher-level causes.




Figure 1: the causal exclusion problem. The arrows of physical causation are bold, the arrows of biological causation not bold.

There are exactly four types of solutions to this problem. One may take supervenience to suggest that all higher-level tokens including biological tokens are *epiphenomena*, thus deleting the arrows starting in  $b_1$  in the figure above. However, in that case, one gets into conflict with well-established results of biology: biological tokens make a difference to what is going on in the world. One may take that making a difference to suggest that biological tokens produce effects that do not have sufficient physical causes, thus proposing a *biophysical interactionism* (downward causation). However, in that case, one contradicts supervenience and completeness, thus provoking a conflict with well-established results of physics.

If one wishes to avoid both these types of conflicts, there are two options. The one is simply to accept all the arrows of causation in the figure above, thus suggesting a *systematic overdetermination*: some biological and some physical effects have, as a matter of law, two sufficient causes, a biological one and a physical one. However, since the biological causes bring about only what is also brought about by physical causes at the same time, there being sufficient physical causes anyway (insofar as there are causes at all), this position leads to the conclusion that biological causes are superfluous, thus coming close to epiphenomenalism. Why should one recognize biological causes if they do not make any difference?

The remaining option is *identity*. If  $b_1$  is identical with  $p_1$  and if  $b_2$  is identical with  $p_2$ , there is no problem of physical causes excluding biological causes, for all biological causes are identical with physical causes. In other words, some physical causes are biological causes. The main argument for the identity of causes is that as soon as one conceives biological causes that are distinct from physical causes, there is a competition between the two types of causes, and the biological causes will inevitably lose that competition, given the completeness of the physical domain. Assuming that property tokens are the relata of causal relations, identity of causes implies token identity: any biological token that is part of a causal chain is identical with a configuration of physical tokens. Since we can suppose that any biological token is integrated into a causal chain, we can say that, for short, any biological token is identical with a configuration of physical tokens.

The causal exclusion problem does not depend on any particular theory of causation – such as, for instance, a theory of causation as the production of effects by causal powers or a theory of causal processes as transfer of a conserved quantity. Even if one settles for a minimal theory of causation in terms of relations of counterfactual dependence between tokens of non-causal, intrinsic properties (Lewis 1973), the brief evaluation of the situation sketched out above does not change. In particular, even in that case, overdetermination is not a plausible way out of the causal exclusion problem (see Sparber 2005 this volume). The second claim mentioned at the beginning of this paper thus is well founded.

However, it seems that the second claim – identity of causes – is in conflict with the first claim, biological kinds as homogeneous kinds in the sense that there is a homogeneous pattern of biological effects that characterizes all and only the tokens that come under a given biological type. The reason is that we cannot conclude from token identity to type identity, that is, the identity of biological kinds with physical kinds. Leaving molecular biology aside, biological kinds typically are functional kinds. Nothing hinders that biological kinds are realized by configurations of physical tokens in such a way that each biological token is identical with a configuration of physical tokens. But, being functional kinds, biological kinds can be multiply realized. That is to say, one and the same biological type  $B$  can be realized by configurations of physical tokens that do not come under a single physical type  $P$ , but under several different physical types  $P_1, P_2, P_3$ , etc. Moreover, this fact of multiple realization is considered to be the main reason for endorsing the first claim. Here is one example taken from Elliott Sober (1993, 76–77):

In addition to explaining single events, we also want to describe *general patterns*. It is here that the vocabulary of supervening properties makes an irreducible contribution. ... For example, the Lotka-Volterra equations in ecology describe how the number of predators and the number of prey organisms are dynamically related. These equations apply to any pair of populations in which organisms in one prey upon organisms in the other. What is this relation of *predation*?

Lions prey on antelopes; Venus's-flytraps prey on flies. What does a lion have in common with a flytrap that makes both of them predators? It isn't in virtue of any physical similarity that these two organisms both count as predators. ... Biological categories allow us to recognize similarities between physically distinct systems. ... It may be that each single event has a physical explanation, but this does not mean that every pattern among events can be characterized in the vocabulary of physics.

If the first claim is based on the assumption that there are biological patterns (biological types or kinds) that cannot be seized by physics and if these patterns are characterized by a homogeneous type of effects that their tokens produce, then how can the second claim be true, namely that each biological cause is identical with a physical cause? Indeed, starting from the second claim, Kim perceives an insoluble conflict and reaches a conclusion that amounts to eliminativism with respect to biological kinds:

... if the "multiplicity" or "diversity" of realizers means anything, it must mean that these realizers are causally and nomologically diverse. Unless two realizers of *E* [emergent property in the sense of higher-level property] show significant causal/nomological diversity, there is no clear reason why we should count them as two, not one. It follows then that multiply realizable properties are ipso facto causally and nomologically heterogeneous. This is especially obvious when one reflects on the causal inheritance principle. All this points to the inescapable conclusion that *E*, because of its causal/nomic heterogeneity, is unfit to figure in laws, and is thereby disqualified as a useful scientific property. ... The conclusion, therefore, has to be this: as a significant scientific property, *E* has been reduced – eliminatively. (Kim 1999, 17–18; see also Kim 2005a, 26, 58, and Kim 2005b this volume, last but one section)

Applying this argument to biological properties in the sense of biological types or kinds, we can sum it up in the following way:

- 1) Each token of a biological type is identical with the physical token that realizes the biological type in the given situation.
- 2) The multiple realizability of biological types implies that the different realizers of a biological type are causally and nomologically heterogeneous.
- 3) Consequently, the biological type is not causally and nomologically homogeneous.

As Kim acknowledges, this conclusion amounts to an eliminative reduction: biological kinds – indeed all higher-level, functional kinds – are eliminated as natural kinds, not figuring in laws. However, this is not what Kim himself wants. In the concluding chapter of his latest book, he clearly states that what he aims at is conservative reduction, not elimination (Kim 2005a, 159–160). The whole point of the argument for biological causes being identical with physical causes is that accepting that identity is the only way to rescue biological causation. However, that argument would lose its point if its consequence were that there are no homogeneous biological kinds. In the following, I shall sketch out two strategies that one can pursue to avoid this eliminativist conclusion – and thus hold on to both the claim of homogeneous biological kinds and the claim of biological causes being identical with physical causes.

## 2. *Two different classifications*

What is the reason for the multiple realization of biological types? Apart from molecular biology, biological types are classified according to their function in the sense of a causal role (pattern of salient characteristic causes and effects). Physical types, by contrast, are classified

according to mereological composition. Classification according to function can differ from classification according to composition in the sense that configurations that are composed in different manners of physical tokens can all exhibit the same salient higher-level effects. That is why biological types can be multiply realized by configurations of physical tokens that come under different physical types due to their different physical composition.

Against this background, we can counter the crucial claim in Kim's argument – namely that (2) the multiple realizability of biological types implies that the different realizers of a biological type are causally and nomologically heterogeneous – in the following way:

1) It is sufficient for the configurations of physical tokens that realize a biological type *B* to be diverse that they are composed in diverse manners of physical tokens. In other words, multiple composition is sufficient for the realizers of a biological type to be diverse, because the realizers are distinguished by their composition.

2) The causal inheritance principle is to say that any token of a biological type *B*, being identical with a configuration of physical tokens, inherits its causal power (that is, its capacity to produce specific effects) from that configuration; for the causal power of that configuration is – or at least includes – the causal power that defines the biological type *B*. That is why the configuration in question is a realizer of *B*. The configuration, in turn, inherits its causal power from – in the last resort – the causal powers of the tokens of fundamental physical properties of which it is composed. The causal power of the configuration supervenes on the causal powers of these tokens.

3) The causal inheritance principle permits that different manners of composition by fundamental physical tokens lead to configurations that produce higher-level effects of the same salient type. No two configurations that are composed in two different ways by fundamental physical tokens produce exactly the same effects. Kim's argument in the citation at the end of the last section is not the trivial point that there is some difference in the effects that different realizers have, but that the different realizers “show significant causal/nomological diversity”. The counter-argument is that the configurations that realize a given biological type, despite being composed in different manners, can have the same pertinent higher-level effects: (a) they all produce effects of the same type, (b) for each of these configurations, these are the significant effects that the configuration produces and (c) only these configurations bring about the effects in question (compare Antony 1999, 19–22, Clapp 2001, 126–132, as well as what Shoemaker 2001, in particular 79, calls forward-looking causal features that are common to all and only the realizers of a given functional type). Consider selection: biological properties of organisms are selected for the effects their instances have. Selection is blind to composition. That is why what is relevant to biology is function instead of composition (Papineau 1993, chapter 2).

In sum, based on the principle of different classifications, there is a way to match the two claims mentioned at the beginning of this paper. Biological causes are identical with physical causes such that any biological token is identical with a configuration of physical tokens. Nothing hinders that configurations that are composed in different manners of physical tokens produce the same salient biological effects – notably effects that are salient for selection in a given environment. Therefore, biological types are causally and nomologically homogeneous; they are heterogeneous only with respect to their physical composition.

This argument vindicates causally homogeneous biological types despite multiple realization and the causal exclusion argument for token identity. Nonetheless, it is

questionable whether this is an argument for the irreducibility of biological theories to theories of chemistry or physics. As quoted in the last section, Sober claims that the vocabulary of supervening properties makes an irreducible contribution to the description of general patterns. However, if we take global supervenience for granted, then any similarity that there is between biological tokens also supervenes on the physical level. In other words, if there is a pattern of similarity among biological tokens that justifies regarding these tokens as belonging to a homogeneous biological type, then there is a supervenience base for that pattern. Consequently, whatever several biological tokens may have in common, this is fixed by what there is on the physical level. But in this case, it seems that what constitutes these patterns – and thus what makes several biological tokens similar – can in principle be accounted for in a physical vocabulary (even if, in practice, our physics may not be in a position to put forward such an account). Hence, in short, the idea of two different manners of classification is sufficient to vindicate the two claims set out at the beginning of this paper, but this idea – and these two claims – do not imply the irreducibility of biological theories.

### 3. *Homogeneous classifications – coarse-grained and fine-grained*

The idea of two different manners of classification pursued in the preceding section is based on the assumption that physics – and chemistry up to molecular biology – classifies configurations of physical tokens according to composition, whereas biology classifies them according to function. Kim's argument ending up with eliminative reduction, by contrast, is based on a causal theory of properties according to which properties are characterized by the effects that their instances characteristically have (Shoemaker 1980). Therefore, it is claimed that if there are different types of realizers of a functional, biological type, there are different types of effects and, consequently, no homogeneous biological type. The weak point of this argument, which was exploited in the last section, is that Kim does not pay heed to the fact that any functional biological type can only be realized by a *configuration* of physical tokens instead of a single physical token – a configuration that can be classified according to composition as well as according to function.

Nonetheless, there is a point about the causal theory of properties. One can claim with good reason that there is in fact no neat distinction between the two manners of classification evoked in the last section. In particular, what counts as a realizer configuration with respect to a higher-level kind may itself be an instance of a functional kind with respect to a lower level, that is, being itself realized by a configuration of lower-level tokens. To take a clear example, assume that mental kinds are functional kinds that are realized by configurations of neurobiological tokens. However, the kinds under which neurobiological tokens come are themselves functional kinds. In the same manner, genes are functional types that are realized by configurations of molecular tokens. One can argue that the molecular kinds in question are themselves functional kinds with respect to more basic chemical or physical kinds (Rosenberg 1994, 24–25).

In this way we get in the end to the idea that there are functional kinds all the way down. There is no need to call into question the assumption that there is a basic level of nature. One can easily define a basic level as the level of properties that need nothing more than a space-time point in order to be instantiated (compare David Lewis' thesis of Humean supervenience in Lewis 1986, IX-X). There is nothing smaller than a space-time point. As regards the fundamental physical kinds, it is undisputed that we can know them only via the relations

their instances enter into (see, for example, Jackson 1998, 23). What is more, there are good physical as well as metaphysical arguments for maintaining that the fundamental physical properties are themselves relational properties, consisting in relations between one space-time point and other space-time points. That is what the position known as structural realism amounts to (e.g. French & Ladyman 2003 and Esfeld & Lam 2005).

This view does not imply that all fundamental properties are causal-functional properties, for there are other relational properties than causal ones. It may even be the case that causal properties do not belong to the fundamental level (see Lam 2005 this volume as regards the physical foundations of causation). Be that as it may, in any case, one can regard all relational properties as functional properties based on a wide notion of functional properties that is not tied to the concept of a causal role, but to the idea of a property being defined by its position in a network with other properties, that is, its relations to other properties. In this case, however, there is no question any more of functional properties being second-order properties that have to be realized by configurations of tokens of first-order properties. There are just tokens of functional-relational properties. Some configurations of tokens of certain lower-level functional-relational kinds are tokens of higher-level functional-relational kinds.

Consequently, there are not two different manners of classification. All classifications proceed in a functional-relational manner. The classifications on the fundamental level are very fine-grained, the classifications on higher-levels are more and more coarse-grained, abstracting from details that are not relevant to the salient effects produced on higher levels. In this perspective, the idea of there being two different kinds of classification is based on a wrong-headed presupposition, namely the view that the fundamental physical properties are intrinsic ones, giving rise to two different manners of classification – according to composition by tokens of intrinsic properties and according to the function that some configurations of tokens of intrinsic properties perform.

Stephen Yablo (1992, 250–257) proposes the following interpretation of the standard view of higher-level kinds being functional kinds whose instances are multiply realized: higher-level functional kinds are determinable kinds, whereas their lower-level realizers are more determinate kinds. Yablo thus integrates the talk of functional roles and their realizers into the familiar talk of – determinable – genus and – more determinate – species. For instance, “animal” is a determinable kind, “mammal” a more determinate kind, and “cat” an even more determinate kind. In the same way, following Yablo, “gene” is a determinable kind, and “certain sequence of bases in the DNA” a more determinate kind.

Yablo (1992, 273–280) sets out the proposition of functional kinds being determinable kinds and their realizers coming under determinate kinds in order to develop a solution to the causal exclusion problem in terms of, in short, causes being proportional to their effects (see Harbecke 2005 this volume). However, in the case of higher-level causation, such as a biological token producing an effect, his proposal amounts to either overdetermination – in case there is a proportional biological and a proportional physical cause of the effect in question – or interactionism, rejecting completeness, in case there is only a biological and not a physical proportional cause for the effect in question. To take an example within the framework of Yablo’s discussion, mental causation, if my right arm goes up this noon because I want to raise it, the effect of my right arm going up is either overdetermined, having a sufficient mental and a sufficient physical (neurobiological) cause, or completeness is violated, if that effect has only a mental cause.

The reason for that impasse is that Yablo conceives determinables and their determinates as really existing, not-identical properties, not-identical types as well as not-identical tokens (Yablo 1992, 266–271; see Gillet & Rives 2005, part 3, for a criticism of this position). However, whereas it is correct to maintain that, for example, a gene of the same type could have been there, even if there had been another sequence of bases in the DNA, one can argue that there is no point in claiming that the same token gene could have been there at the same instant, even if there had been another sequence of bases in the DNA at that time (see Sparber 2005 this volume) – in the same way as it would be absurd to claim that the same blue token could have been there, even if there had been another shade of blue (say aquamarine instead of cobalt). In those cases, there would have been another token, albeit of the same type – that is, another token gene or another token blue. Hence, pace Yablo, the argument for regarding higher-level causes, such as biological ones, as being identical with physical causes stands. Consequently, higher-level tokens, such as biological ones, are identical with configurations of physical tokens.

Token identity implies that it is the same individual that makes true a physical description as well as higher-level descriptions such as a biological one. In short, there are different concepts and different degrees of abstract descriptions – from a detailed, determinate description in terms of fundamental physics to more determinate descriptions in a chemical, a molecular biological, or a genetic vocabulary, etc. – but all these descriptions refer to the same individuals and have the same truthmakers (compare Heil 2003, chapters 2 to 7).

Leaving aside Yablo's ontology, we can employ his suggestion about functional kinds being determinable kinds to spell out the idea of a homogeneous classification – and, pace Kim, we get homogeneous biological kinds, this time without relying on two different manners of classification. A detailed description in terms of fundamental physical concepts is a very determinate description, higher-level descriptions are less and less determinate. Nonetheless, they refer to homogeneous patterns of characteristic effects, and their concepts figure in laws. Determinable kinds are homogeneous kinds, being defined notably by a homogeneous pattern of effects. They are just less determinate than their determinates. A genus is something homogeneous (as the word implies), as are its species. For instance, the genus “mammal” is characterized by the effect of giving birth to live babies and feeding them on milk from the breast – that is the salient feature that all and only the tokens of mammals have in common.

If the relation between functional properties and their realizers is the one of determinables to their determinates, there is nothing particular about multiple realization. The same gene is realized by different chunks of DNA as mammals are realized by cats, dogs, whales, elephants, etc. – if one wishes to employ the term “realization” in that case. Of course, what one says about a gene is true not only of a chunk of DNA of a particular composition, but also true of chunks of DNA of other composition – in the same way as what one says about mammals is true not only of cats, but also of whales. However, this triviality does not hinder theory reduction.

As is clear since the improvement by Hooker (1981) on the account of theory reduction of Nagel (1961, chapter 11), in order to reduce a higher-level theory to a lower-level theory, it is necessary to prepare the lower-level theory for reduction: one has to introduce within the detailed descriptions of the lower-level theory more abstract concepts tailor-made for the specific configurations of lower-level tokens to which a higher-level theory refers (see in

particular Hooker 1981, 49 and 498–508). If all properties are functional-relational properties, there is no principled problem for introducing such concepts in a way that is suitable for theory reduction (see Sachse 2005 this volume). In sum, in whatever way one cuts the dough, one gets to reduction, but to conservative reduction – reduction without elimination, respecting both the identity of biological causes and tokens with configurations of physical causes and tokens as well as the homogeneity of biological kinds, characterized each by a homogeneous pattern of salient effects that their instances produce.

### References

- Antony, L. M., 1999, “Multiple realizability, projectibility, and the reality of mental properties”, *Philosophical Topics*, 26: 1–24.
- Clapp, L., 2001, “Disjunctive properties: multiple realizations”, *Journal of Philosophy*, 98: 111–136.
- Esfeld, M. & Lam, V., 2005, “Moderate structural realism about space-time”, Manuscript, submitted.
- French, S. & Ladyman, J., 2003, “Remodelling structural realism: Quantum physics and the metaphysics of structure”, *Synthese*, 136: 31–56.
- Gillet, C. & Rives, B., 2005, “The non-existence of determinables: or, a world of absolute determinates as default hypothesis”, *Noûs*, 39: 483–504.
- Harbecke, J., 2005, “A realist approach to biological events”, this volume.
- Heil, J., 2003, *From an ontological point of view*, Oxford: Oxford University Press.
- Hooker, C. A., 1981, “Towards a general theory of reduction. Part I: Historical and scientific setting. Part II: Identity in reduction. Part III: Cross-categorical reduction”, *Dialogue*, 20: 38–60, 201–236, 496–529.
- Jackson, F., 1998, *From metaphysics to ethics. A defence of conceptual analysis*, Oxford: Oxford University Press.
- Kim, J., 1998, *Mind in a physical world. An essay on the mind-body problem and mental causation*, Cambridge (Massachusetts): MIT Press.
- Kim, J., 1999, “Making sense of emergence”, *Philosophical Studies*, 95: 3–36.
- Kim, J., 2005a, *Physicalism, or something near enough*, Princeton: Princeton University Press.
- Kim, J., 2005b, “Laws, causation, and explanation in the special sciences”, this volume.
- Lam, V., 2005, “Causation and space-time”, this volume.
- Lewis, D., 1973, “Causation”. *Journal of Philosophy*, 70: 556–567.
- Lewis, D., 1986, *Philosophical papers. Volume 2*, Oxford: Oxford University Press.
- Nagel, E., 1961, *The structure of science. Problems in the logic of scientific explanation*, London: Routledge.
- Papineau, D., 1993, *Philosophical naturalism*, Oxford: Blackwell.
- Papineau, D., 2002, *Thinking about consciousness*, Oxford: Oxford University Press.
- Rosenberg, A., 1978, “The supervenience of biological concepts”, *Philosophy of Science*, 45: 368–386.
- Rosenberg, A., 1994, *Instrumental biology or the disunity of science*, Chicago: University of Chicago Press.
- Sachse, C., 2005, “Reduction of biological properties by means of functional sub-types”, this volume.
- Shoemaker, S., 1980, “Causality and properties”. In: van Inwagen, P. (ed.), *Time and cause*, Dordrecht: Reidel, 109–135. Reprinted in Shoemaker, S., 1984, *Identity, cause, and mind. Philosophical essays*, Cambridge: Cambridge University Press, essay 10.
- Shoemaker, S., 2001, “Realization and mental causation”. In: Gillet, C. & Loewer, B. (eds.), *Physicalism and its discontents*, Cambridge: Cambridge University Press, 74–98.
- Sober, E., 1993, *Philosophy of biology*, Boulder: Westview Press.
- Sparber, G., 2005, “Counterfactual overdetermination vs. the causal exclusion problem”, this volume.
- Yablo, S., 1992, “Mental causation”, *Philosophical Review*, 101: 245–280.