Functional Stability and Systems Level Causation

Anders Strand and Gry Oftedal†‡

A wide range of gene knockout experiments shows that functional stability is an important feature of biological systems. On this backdrop, we present an argument for higher-level causation based on counterfactual dependence. Furthermore, we sketch a metaphysical picture providing resources to explain the metaphysical nature of functional stability, higher-level causation, and the relevant notion of levels. Our account aims to clarify the role empirical results and philosophical assumptions should play in debates about reductionism and higher-level causation. It thereby contributes to the development of a philosophical foundation for systems biology.

1. Introduction. Issues in metaphysics are often discussed in isolation from empirical results provided by special sciences like biology. In the same way, theoretical issues in biology are often discussed without sufficient sensitivity to the metaphysical assumptions and commitments they involve. The prospects for new insights and advancements by bringing contemporary metaphysics and biology closer together are therefore promising.

Systems biology is a relatively new but rapidly expanding discipline. Nonreductive assumptions often associated with this approach have fueled the interest in the philosophical aspects of systems biology. Systems biologists aim to understand how a wide range of functional properties of living organisms is brought about by systems behavior and by the interactions of system constituents (Alberghina and Westerhoff 2005). Biological laboratory research, mathematical modeling, complex statistics, and biophysical methods are all in play in systems biology research. Although there are both "bottom-up," "top-down," and "middle-out" ap-

†To contact the authors, please write to: CSMN/IFIKK, University of Oslo, Box 1020 Blindern, N-0315, Oslo, Norway; e-mail: anders.strand@ifikk.uio.no, gry.oftedal@ifikk.uio.no.

‡Thanks to Henrik Forssell, Nils Roll-Hansen, and the participants at the biennial PSA meeting in Pittsburgh 2008 and the Seminar in Science Studies in Oslo, January 2009.

Philosophy of Science, 76 (December 2009) pp. 809–820. 0031-8248/2009/7605-0024$10.00
Copyright 2009 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/605811 Published online by Cambridge University Press

proaches, the discipline is often described as shifting the focus in biology from the level of single molecules to system levels (Boogerd et al. 2007, 4).

On the other hand, reduction and the status of higher-level causes are both central and recurrent topics in contemporary philosophy of science. There is broad consensus that higher-level explanations are heuristically useful and that higher-level explanations and theories do not translate directly into lower-level explanations and theories. However, discussions concerning the *possibility of* and *prospects for* reductions have flourished for several decades (e.g., Nagel [1961] 1979; Bickle 2003). There is also a widespread causal exclusion worry that higher-level entities *must be* reducible to lower-level entities in order to be causally efficacious (Kim 1998). It is pertinent to ask how these topics relate to the philosophical understanding of systems and of systems biology.

Biological systems are often individuated relationally, that is, in relation to their environment, their functions, and their history. Such systems are, for instance, cells, tissues, whole organisms, or ecosystems. Questions then arise about the relation between such relationally individuated systems and their specific structural implementation. In this respect, an important feature of biological systems is *functional stability*. Biological systems often retain some or all of their functions under different alterations of their structure. This feature, we argue, should influence how we build a philosophical account of the relation between systems and their structural implementation. Consider a system that is in part functionally individuated. Such a system can be modally robust with respect to its structural implementation. In other words, such a system might survive structural alterations that preserve its functions. This feature, we suggest, is sufficient for denying identity between the system and its structural implementation. We argue that an understanding of levels in terms of a *realization relation* is more fruitful.

Functional stability is arguably a widespread property of biological systems (e.g., Kitano 2004; Wagner 2005) and is strikingly exemplified by a range of gene knockout experiments (e.g., Shastry 1994). The outcomes of these experiments provide an empirical backdrop for our account of the relation between systems and their constituents. The concepts of levels, realization, and nonextensional mereology will here prove to be of great importance.[1] Since a significant goal is to develop a philosophical account of the relation between scientific levels, we find it important to keep the questions of reductionism and higher-level causation open at the outset. The account we suggest will in conclusion lend plausibility to an argument

1. The term "mereology" refers to formal theories of the part-whole relation.

https://doi.org/10.1086/605811 Published online by Cambridge University Press

for higher-level causation based on counterfactual dependence, but it leaves the argument susceptible to empirically based criticisms.

We proceed in three main steps. First, we present empirical support for functional stability and formulate a more precise notion of functional stability. Second, we rehearse a counterfactual dependence-based argument for systems level causation. Third, we sketch a metaphysical picture that facilitates the philosophical understanding of these issues, while still leaving the specific conclusions sensitive to empirical constraints.

2. Gene Knockouts. The past 25 years have seen gene knockouts become an increasingly more important experimental tool in molecular biology. The basic idea is to inactivate one or several genes in order to reveal the connections between those genes and their hypothesized functions. Constructed DNA is substituted for some of the genetic material of individual cells, for instance, in early embryos of mice. The resulting organisms are able to pass on the disrupted gene to future generations. M. R. Capeschi, M. J. Evans, and O. Smithies won the 2007 Nobel Prize in Physiology and Medicine for the development of gene knockout techniques.

Many gene knockout experiments involving assumed functionally important genes have produced rather unexpected results (Shastry 1994; Sampogna and Nigam 2004; Schuendeln et al. 2004). Often organisms with an inactivated gene develop into well-functioning phenotypes with no abnormalities. In yeast, more than 40% of the genes have been knocked out without any detectable effect on the phenotype in five different environments (Smith et al. 1996; Wagner 2000). Numerous genes are also dispensable in mice (Shastry 1994). Additionally, there are many cases in which the phenotype changed less or in a different way than researchers expected, given the assumed function of the gene (see reviews by Shastry [1994] and Sampogna and Nigam [2004]). Many genes show functional redundancy insofar as they are replaceable by other genes or mechanisms in the organism without significantly affecting (if at all) higher-level causal properties.

One example is the RARα1 gene, which contributes to the production of a retinoic acid receptor. Retinoic acid is a vitamin A residue and strongly involved in embryo development. One would therefore expect mice without the RARα1 to show embryonic defects. However, the mice carrying the disrupted RARα1 showed no abnormalities or changes during development (Li et al. 1993). Suggested explanations are that receptors are produced by a backup mechanism or that RARα1 is simply not involved in embryonic development (Li et al. 1993). It seems that the lack of abnormalities in the case of the knocked out RARα1 can be explained by functional redundancy in the sense that receptors were produced through alternative pathways. The fact that severe abnormalities appeared

https://doi.org/10.1086/605811 Published online by Cambridge University Press

when disrupting the whole RAR cluster of genes, not just the RARα1, adds support to this explanation. Embryonic development is in this case not counterfactually dependent on the RARα1 gene but, rather, on the RAR cluster or network of genes and regulatory mechanisms.

Similar backup mechanisms are activated for hox genes. Hox gene clusters are regarded as being of particular importance to development, and their structure, as well as their contribution to mechanisms and functions, is highly conserved across evolution and time. Nonetheless, a knockout of hox-3.1 did not produce any expected changes in the phenotype (LeMouellic, Lallemand, and Brulet 1992). And, although the knockouts of hox-1.5 and hox-1.6 did produce some deficiencies, they did not result in the expected major disruptions of development (Chisaka and Capecchi 1991; Chisaka, Musci, and Capecchi 1992).

Thus, a gene is dispensable if backup genes can take its place when it is knocked out. There is also a second type of stability (or robustness) of biological systems that does not depend on duplicate genes. Many genes without backup copies in the genome can also be deleted without having any phenotypic effects. According to, for example, Wagner (2005), so-called *distributed robustness* is widespread and relies on many different parts contributing to system functions. A system can compensate for a missing part through network adjustments involving many parts, not only one or a few genes. System stability may not result from gene sequence similarity, as is often suggested, but, instead, from interactions between unrelated genes in a manner that can be described as a systemic property (Wagner 2000, 2005). In these cases, the counterfactual dependence of the relevant effect shifts to a higher system level rather than to clusters of genes as in the RARα1 case.

The concept of functional stability is in use in several biological disciplines and is applied to systems in, for example, ecological research, where the functional stability of soil systems is found to decrease when biodiversity is lowered (e.g., Griffiths et al. 2001). In this context, functional stability is measured in terms of system resistance (the ability of a system to withstand the immediate effects of perturbation) and system resilience (the ability of a system to recover from perturbation).

In network theory, we find a similar concept of functional stability (Kauffman 1967; Albert, Jeong, and Barabasi 2000). Sampogna and Nigam (2004) suggest that a system level approach to network structures may be the best explanatory framework for the lack of phenotypic effects of gene knockouts. Some architectural features have been found to be common to several types of complex systems resulting in robust and relatively error-tolerant networks (Albert et al. 2000). Complex networks in organisms seem to consist of a collection of interconnected nodes (e.g., genes or other molecules) where the level of connectivity is unevenly dis-

https://doi.org/10.1086/605811 Published online by Cambridge University Press

tributed. Some nodes exhibit many more links than others in that they interact with a high number of other nodes. However, nodes of low connectivity predominate. The highly connected nodes are apparently significantly more important than the others in shaping the processes of the network. Statistically, errors most often occur in the less significant nodes (since they are more frequent), and these perturbations will not have critical consequences for the functioning of the system.

Based on gene knockout experiments, ecology, and network theories, we formulate the following definition of functional stability:

> **Functional Stability.** The ability of a system to retain its higher-level functional properties when parts of the system are subjected to perturbation. Perturbation may take the form of removing, destroying, or replacing parts of the system.

Our formulation of functional stability resembles the concept of robustness in biology. For instance, Wagner (2005, 176) holds that "a biological system is robust to mutations if it continues to function after genetic changes in its parts."

Functional stability, here exemplified by gene knockout experiments, definitely feeds interesting empirical evidence into the philosophical debate on levels, causal exclusion, and higher-level causation. The fact that higher-level properties on several occasions are found to be independent of lower-level changes shows differences in the patterns of counterfactual dependence between an effect and lower- and higher-level cause candidates. This feature can be used in an argument for higher-level causes, to which we now turn.

3. Systems Level Causation and Counterfactual Dependence. In this section we present an argument for systems level causes in biology. It requires a metaphysical framework within which the difference between—and the status of—different cause candidates can be precisely understood and explained. We discuss this framework in the last section. The argument basically involves an empirical comparison of the causal status of systems level causes and their specific structural implementation. We achieve this by comparing the patterns of counterfactual dependence between higher and lower-level cause candidates and their alleged effects. As we saw in the previous section, these patterns differ between the levels. Insofar as counterfactual dependence is symptomatic of causation, the differences can be marshaled in support of higher-level causation. Similar arguments have been discussed in the field of mental causation in philosophy of mind (e.g., Yablo 1992).

Causal relations are generally thought to be supported by counterfactuals. It is arguably an inherent assumption of the methodology of the

https://doi.org/10.1086/605811 Published online by Cambridge University Press

gene knockout experiments as well, since the basic idea is to inactivate a gene and see what difference it makes to the phenotype. In other words, one investigates whether the effect still happens in the absence of the alleged cause. If the effect does not occur, there is reason to believe the alleged cause to be an actual cause of the phenotypic effect in question.

In the examples discussed above, however, the phenotypic effect does not depend counterfactually on the presence of the knocked out gene, and thus not on this *specific* structural implementation of the system. Prima facie, this fact leaves us with two options: (i) the knocked out gene is not a cause after all, or (ii) the case is akin to so-called *preemption cases* known from the causation literature. A central feature of preemption cases is the absence of counterfactual dependence between cause and effect. They have been offered as counterexamples to the counterfactual analysis of causation and sparked much discussion (e.g., Lewis 1986; Collins, Hall, and Paul 2004). According to ii, the knocked out gene is a preemptive cause; it is responsible for the effect; *and* it preempts a backup mechanism from causing the effect. When the relevant gene is knocked out, the backup mechanism is no longer preempted, and the effect still occurs. Even on this interpretation distributed backup mechanisms indicate the reality of systems level causes.

However, if such examples are considered on the assumption that the system itself is a genuinely different entity from its specific structural implementation, a third option becomes available. At the level of the more modally robust system there will be a cause candidate C_1. It will be realized by a modally fragile lower-level cause candidate C_2, which is identical to the sum of all the molecules and ions that realize the system.[2] Recall the RARα1 example. From a higher-level perspective, the RARα1 gene is part of the system causing normal embryonic development. However, RARα1 is not an *essential part* of this system, since its inactivation does not affect embryonic development. Candidate C_2 shares all proper parts with C_1 at a certain time but differs from C_1 in that C_2 has all its parts essentially. This explains the different patterns of counterfactual dependence between the effect and C_1 and the effect and C_2; the effect occurs in worlds where C_2 does not occur, but the effect does not occur in worlds where C_1 does not occur. As we interpret the empirical results

2. The relations between the constitutive molecules are important for the functioning of the whole. Questions could be raised concerning the sum's (C_2) sensitivity to alterations in these relations and whether it is therefore modally robust in terms of organization. This needs to be investigated. Also, the sum of microparts has all of its microparts essentially, although we do not assume mereological essentialism (the claim that all parts are essential to a whole) as a general thesis.

https://doi.org/10.1086/605811 Published online by Cambridge University Press

(Li et al. 1993), normal embryonic development depends counterfactually on C_1 and not on C_2.[3]

This argument consists of a schema for empirically establishing a nonreductive account of higher-level causation and of empirical results that feed into the schema in the right way. However, the schema should not be question begging, and it should allow for the crucial assumption that systems are—or at least can be—genuine entities apart from their specific structural implementation. We therefore turn to a metaphysical account of realization and levels, with a view to justify and explain this part of the argument.

4. Metaphysics of Levels and Causal Relata. In spite of important contributions (e.g., Oppenheim and Putnam 1958; Schaffer 2003; Craver 2007), the notion of levels is typically a tacit backdrop against which other questions in philosophy are being discussed. Arguably, the concept has not been the focus of sufficient critical discussion. What Jaegwon Kim says is representative. Kim speaks of

> the familiar multilayered model that views the world as stratified into different "levels", "orders", or "tiers" organized in a hierarchical structure. The bottom level is usually thought to consist of elementary particles, or whatever our best physics is going to tell us are the basic bits of matter out of which all material things are composed. As we go up the ladder, we successively encounter atoms, molecules, cells, larger living organisms, and so on. The ordering relation that generates the hierarchical structure is the mereological (part-whole) relation: entities belonging to a given level, except those at the very bottom, have an exhaustive decomposition, without remainder, into entities belonging to the lower levels. (Kim 1998, 15)

Let's call this the mereological conception of levels. It accords with Oppenheim and Putnam's seminal paper, "The Unity of Science as a Working Hypothesis" (1958), which has shaped discussions of reduction and scientific unity. By contrast, the level hierarchy we are primarily interested in here is generated by a realization relation and does not track the mereological level hierarchy.

Let us first consider some intuitive and some not-so-intuitive examples of realization:

1. His dream guitar was realized in a custom ordered gold sparkling Stratocaster.

3. Empirical results could favor intermediate cause candidates between the two candidates we are discussing here.

https://doi.org/10.1086/605811 Published online by Cambridge University Press

2. Redness is realized in the car outside my window.
3. A cell at time t is realized by the totality of its cell components at time t.

In everyday language a realization relation is usually said to hold between an abstract idea and a concrete object, as in 1. Sometimes it is thought to obtain between a type and a particular token, as in 2. However, we are interested in cases in which one concrete object realizes another concrete object, as in 3, where the cell can be seen as one object and the totality of cell constituents as another object that realizes the cell at time t. In such cases, we claim, the realized object and its realizer are different objects but which nevertheless *share all proper parts*. We also suggest that the realization relation can be defined in terms of how the objects relate to their shared parts.

The term *mereology* refers to theories of the part-whole relation (see Simons 1987 and Varzi 2006 for background and discussions). According to a so-called extensional closure mereology (CEM), the part-whole relation is a partial ordering supplemented by a supplementation principle and two closure principles (sum and product).[4] According to CEM the part-whole relation has the following properties:

1. Everything is a part of itself (reflexive).
2. If A is a part of B and B is a part of A, then $A = B$ (antisymmetric).
3. If A is a part of B and B is a part of C, then A is a part of C (transitive).

The first three properties define a partial ordering.[5]

4. If A is not a part of B, then there exists a C that is part of A and that does not overlap with B (strong supplementation).
5. If A and B underlap, then there exists a C such that for all D: C overlaps D if and only if D overlaps A or D overlaps B (sum).
6. If A and B overlap, then there exists a C such that for all D: D is a part of C if and only if D is a part of A and D is a part of B (product).

4. The presentation is based on Varzi 2006. Overlap and underlap appear as undefined terms in the axioms in the main text. Two objects *overlap* if they have a part in common, and two objects *underlap* if they are both parts of the same object. Simons (1987) formulates the mereological theory in terms of *proper* parthood, but that is only a matter of convenience. A is a proper part of B if A is a part of B and $A \neq B$.

5. Note that a partial ordering on a set allows for pairs $\langle a, b \rangle$ that are not ordered. The mereological conception of levels will therefore be incomplete in the sense that it does not relate entities that are not related as part and whole. For example, it provides no answer to questions like: "Is his new gold sparkling guitar at a lower level than the Humboldt University?"

https://doi.org/10.1086/605811 Published online by Cambridge University Press

	$\mathbf{P_1}$	$\mathbf{P_2}$	$\mathbf{P_3}$
$\mathbf{O_1}$	E	E	E
$\mathbf{O_2}$	E	E	A
$\mathbf{O_3}$	E	A	E
$\mathbf{O_4}$	E	A	A
$\mathbf{O_5}$	A	A	A
$\mathbf{O_6}$	A	A	E
$\mathbf{O_7}$	A	E	A
$\mathbf{O_8}$	A	E	E

Figure 1.

A theorem of CEM is this: if two objects have exactly the same proper parts, they are identical. In order to spell out the nature of a nontrivial realization relation obtaining between objects that share all proper parts, one has to deny this theorem. This can be done by adopting a nonextensional mereological framework. A nonextensional mereology can be developed in different ways; we suggest denying the strong supplementation principle and adopting a distinction between *essential* and *accidental* parthood. Any part of a whole is either an essential or an accidental part. A proper part is essential to a whole if the whole of which it is part cannot exist without that part. A proper part is accidental to a whole if the whole of which it is part can exist without it. The development of a suitable nonextensional framework is a topic for further work.[6]

We spell out the realization relation we have in mind in terms of how a realized entity and its realizer relate to their shared parts. A simple model illustrates the idea. It consists of three atoms P_1, P_2, and P_3 and all of the objects composed of these three atoms. We also assume that a universal sum axiom holds, saying that any collection of parts constitutes an object. These assumptions secure the existence of objects $O_1, \ldots, O_8$ composed of P_1, P_2, and P_3. In Figure 1, E means that the object O_n has the part P_m essentially, while A means that it has it accidentally.

If two objects differ in the parts they have essentially and accidentally, they cannot be identical. We suggest that an object either has all its parts essentially or is realized by an object that has all its parts essentially. For

6. Simons (1987) discusses nonextensional mereologies and modal mereologies, and his work would be a natural point of departure for such investigations. We allow ourselves to talk in terms of a distinction between essential and accidental parts in the main text above.

https://doi.org/10.1086/605811 Published online by Cambridge University Press

example, O_1 is a realizer of O_2, . . . , O_8. Moreover, O_2 and O_3 have a common realizer, but neither realizes the other; O_4 is realized by O_1, O_2, and O_3. In general:

> **Realization among Objects.** Object O is realized by object R if and only if O and R share all proper parts, and O is strictly more modally robust than R with respect to their proper parts.

Object O is strictly more modally robust than R with respect to their proper parts if and only if all proper parts essential to O are essential to R and R has at least some proper parts essentially that O has accidentally.

An organism, for example, is realized by the totality of its constitutive ions and molecules at any one time. This totality, however, cannot survive replacement of, for example, individual molecules; if any molecule is replaced, a different totality results. Such totalities have all their parts *essentially*; they are defined as that whole precisely by having those parts. The organism, on the other hand, does survive the replacement of (at least some) individual molecules.

As shown by the gene knockout examples, a system level cause for a specific phenotype can survive the replacement of assumed functionally important genes. Accordingly, the system level cause is more modally robust with respect to its proper parts than its realizer. On the realization relation we favor, the system level cause will be realized by—but not be identical to—its specific structural implementation. The same relation obtains between biological systems in general and the totalities of their constituents, we believe.

If biological systems do turn out to relate to their realizers by the kind of realization relation defined above, it becomes evident why the argument for higher-level causation has plausibility. The presence profile of higher-level entities across possible worlds differs from the presence profile of their modally fragile realizers. Thus, it should not be surprising that the patterns of counterfactual dependence between an effect and different cause candidates, related by the realization relation, differ.

A modally robust entity, like a system with a certain degree of functional stability, will reasonably be present in more possible worlds than any of its specific implementations.[7] When considering causal relations, as exemplified by the gene knockout experiments, one may find that these differences in presence profiles favor the causal status of the higher-level entities. This also shows the importance of empirical results for the debate about reductionism.

7. There are complications here, since a system might be less modally robust than the sum of molecules and ions that realize it in terms of organization.

https://doi.org/10.1086/605811 Published online by Cambridge University Press

5. Concluding Remarks. We have described a metaphysical framework important for questions about higher-level causation. If we are right, in order not to be question begging, the debate itself is committed to the view that distinct objects at different levels can share all proper parts, as we suggest biological systems and their structural realizers actually do. The results of the gene knockout experiments reported above empirically support that higher-level entities sometimes have causal priority over their lower-level realizers.

On a more general note, a philosophical foundation for systems biology is needed as the discipline of systems approaches currently expands and is by many expected to take over the scene of cutting-edge biology. We have contributed some steps toward this goal by suggesting a metaphysical background picture for the nature of functional stability of living systems and by showing how the causal efficacy of entities at the system level can be established.

As we have aimed to show through the course of this paper, new directions in biology can give valuable information to established philosophical discussions. We expect molecular and systems biology, as well as the presently suggested metaphysical framework, to shed further light on the issues of reductionism and causation, an area of study to which we hope to contribute.

REFERENCES

Alberghina, L., and H. V. Westerhoff, eds. (2005), *Systems Biology: Definitions and Perspectives*. Berlin and Heidelberg: Springer-Verlag.

Albert, R., H. Jeong, and A.-L. Barabasi (2000), "Error and Attack Tolerance of Complex Networks", *Nature* 406: 378–382.

Bickle, John (2003), *Philosophy and Neuroscience: A Ruthlessly Reductive Account*. Dordrecht: Kluwer.

Boogerd, Fred C., Frank J. Bruggeman, Jan-Hendrik S. Hofmeyr, and Hans V. Westerhoff (2007), "Towards Philosophical Foundations of Systems Biology: Introduction", in Fred C. Boogerd, Frank J. Bruggeman, Jan-Hendrik S. Hofmeyr, and Hans V. Westerhoff (eds.), *Systems Biology: Philosophical Foundation*. Amsterdam and Oxford: Elsevier, 1–19.

Chisaka, O., and M. R. Capecchi (1991), "Regionally Restricted Developmental Defects Resulting from Targeted Disruption of the Mouse Homeobox Gene hox-1.5", *Nature* 350: 473–479.

Chisaka, O., T. S. Musci, and M. R. Capecchi (1992), "Developmental Defects of the Ear, Cranial Nerves and Hindbrain Resulting from Targeted Disruption of the Mouse Homeobox Gene hox-1.6", *Nature* 355: 516–520.

Collins, John, Ned Hall, and L. A. Paul, eds. (2004), *Causation and Counterfactuals*. Cambridge, MA: MIT Press.

Craver, Carl (2007), *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Clarendon.

Griffiths, B. S., M. Bonkowski, J. Roy, and K. Ritz (2001), "Functional Stability, Substrate Utilisation and Biological Indicators of Soils Following Environmental Impacts", *Applied Soil Ecology* 16: 49–61.

https://doi.org/10.1086/605811 Published online by Cambridge University Press

Kauffman, S. A. (1967), "Metabolic Stability and Epigenesis in Randomly Constructed Genetic Nets", *Journal of Theoretical Biology* 22: 437–467.
Kim, Jaegwon (1998), *Mind in a Physical World*. Cambridge, MA: MIT Press.
Kitano, H. (2004), "Biological Robustness", *Nature Reviews Genetics* 5: 826–837.
LeMouellic, H., Y. Lallemand, and P. Brulet (1992), "Homeosis in the Mouse Induced by a Null-Mutation in the hox-3.1 Gene", *Cell* 69: 251–264.
Lewis, David (1986), "Causation", in *Philosophical Papers*. Vol. 2. New York and Oxford: Oxford University Press, 159–213.
Li, E., H. M. Sucov, K. F. Lee, R. M. Evans, and R. Jaenisch (1993), "Normal Development and Growth of Mice Carrying a Targeted Disruption of the α-1 Retinoic Acid Receptor Gene", *Proceedings of the National Academy of Sciences of the United States of America* 90: 1590–1594.
Nagel, Ernest ([1961] 1979), *The Structure of Science: Problems in the Logic of Scientific Explanation*. Indianapolis and Cambridge: Hackett.
Oppenheim, Paul, and Hilary Putnam (1958), "The Unity of Science as a Working Hypothesis", in Herbert Feigl, Grover Maxwell, and Michael Scriven (eds.), *Minnesota Studies in the Philosophy of Science*, vol. 3, *Concepts, Theories, and the Mind-Body Problem*. Minneapolis: University of Minnesota Press, 3–36.
Sampogna, R. V., and S. K. Nigam (2004), "Implications of Gene Networks for Understanding Resilience and Vulnerability in the Kidney Branching", *Physiology* 19: 339–347.
Schaffer, Jonathan (2003), "Is There a Fundamental Level?", *Noûs* 37: 498–517.
Schuendeln, M. M., R. P. Piekorz, C. Wichmann, Y. Lee, P. J. McKinnon, K. Boyd, Y. Takahashi, and J. N. Ihle (2004), "The Centrosomal, Putative Tumor Suppression Protein TACC2 Is Dispensable for Normal Development, and Deficiency Does Not Lead to Cancer", *Molecular and Cellular Biology* 24: 6403–6409.
Shastry, B. S. (1994), "More to Learn from Gene Knockouts", *Molecular and Cellular Biochemistry* 136: 171–182.
Simons, Peter (1987), *Parts: A Study in Ontology*. Oxford: Oxford University Press.
Smith, V., K. N. Chou, D. Lashkari, D. Botstein, and P. O. Brown (1996), "Functional Analysis of the Genes of Yeast Chromosome V by Genetic Footprinting", *Science* 274: 2069–2074.
Varzi, Achille (2006), "Mereology", in Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Winter 2006 ed. http://plato.stanford.edu/archives/win2006/entries/mereology/.
Wagner, Andreas (2000), "Robustness against Genetic Mutations in Genetic Networks of Yeast", *Nature Genetics* 24: 355–361.
——— (2005), "Distributed Robustness versus Redundancy as Causes of Mutational Robustness", *Bioessays* 27: 176–188.
Yablo, Stephen (1992), "Mental Causation", *Philosophical Review* 101: 245–280.

https://doi.org/10.1086/605811 Published online by Cambridge University Press