

# Which Kind of Causal Specificity Matters Biologically?

Marcel Weber\*†

---

Paul Griffiths et al. have proposed a quantitative measure of causal specificity and used it to assess various attempts to single out genetic causes as being causally more specific than other cellular mechanisms, for example, alternative splicing. Focusing in particular on developmental processes, they have identified a number of important challenges for this project. In this discussion note, I would like to show how these challenges can be met.

---

**1. Introduction.** The concept of causal specificity is of considerable interest to the philosophy of biology because it promises to provide a rationale for the appeal of a certain kind of explanatory strategy that has been very prominent in twentieth- and twenty-first-century biology. This strategy consists in focusing explanations on biological entities that are often described as information bearing, a kind of description that many take to be metaphorical (e.g., Sarkar 1996; Griffiths 2001).<sup>1</sup> The usual suspects for such biological entities include genomic DNA as well as messenger-RNA (mRNA), which feature prominently in countless biological explanations and are often described as “determining” the sequences within proteins or RNAs. However, as proponents of Developmental Systems Theory have pointed out, we should re-

Received January 2016; revised August 2016.

\*To contact the author, please write to: Department of Philosophy, University of Geneva, 5 rue De-Candolle, 1211 Genève, Switzerland; e-mail: marcel.weber@unige.ch.

†I wish to thank Lorenzo Casini, Michal Hladky, Maximilian Huber, Alan Love, Guillaume Schlaepfer, Christian Sachse, Ken Waters, Bill Wimsatt, and the anonymous reviewers for helpful criticism on various drafts. Versions of this article were presented at the University of Bristol (Conference on Causation in Biology, September 2015) and the University of Bordeaux/CNRS (Philosophy and Biology Seminar, April 2016), and it has greatly benefited from discussions there.

1. Griffiths and Stotz (2013, chap. 6) argue that the notion of genetic information is metaphorical but that such metaphors play a constructive role in science.

Philosophy of Science, 84 (July 2017) pp. 574–585. 0031-8248/2017/8403-0009\$10.00  
Copyright 2017 by the Philosophy of Science Association. All rights reserved.

spect a principle of “causal parity” or “parity of reasoning” and not take it for granted that such entities actually do play such a special causal role (e.g., Oyama 2000; Griffiths and Gray 2005). The attribution of any such role requires some positive account of what exactly that special role consists in, which has proved to be more difficult than one might have thought.

A promising attempt to provide such an account has used the idea that what distinguishes genetic material from other biological causes is their *causal specificity* (Waters 2007; Woodward 2010; for a critique, see Griffiths and Stotz 2013, chap. 4).<sup>2</sup> According to this idea, some causes allow a much more fine-grained control over their effect variable than others. For example, a light dimmer allows a more fine-grained control of light intensity than a simple toggle switch. Similarly, so it is argued, DNA and RNA sequence variation is a causal difference maker that has a more fine-grained control over their effects than other parts of a living cell, for example, the enzymes that are necessary for protein synthesis. Is it this fine-grain or specificity that justifies the biologists’ highlighting of certain biomolecules as information bearers, even if the latter description should still be metaphorical?

Existing attempts to answer this question in the affirmative are largely qualitative. Waters (2007) as well as Woodward (2010) treated causal specificity as a property of causal relations that is either present or absent, even though they both hinted that it admits of degrees. In Weber (2006, 2013), I considered causal specificity to be a matter of degrees and argued that some genetic causes, namely, DNA and mRNA, manifest a greater degree of specificity with respect to RNA or protein sequence than other causal factors involved in gene expression (Waters [2007] tentatively accepts this claim in a footnote). However, no quantitative measure of the degree of causal specificity was available then. Thus, Griffiths et al. have done us a great service in developing a way of measuring it. In fact, they provide the first precise definition of this notion. Basically, they identify causal specificity with the mutual information that interventions on a cause variable and the value of an effect variable bear on each other. Thus, basic concepts of information theory turn out to be helpful.

Using this information-theoretic specificity measure, Griffiths et al. venture to compare different causes of organismic development with respect to their causal specificity. They consider three kinds of causal specificity, namely, the specificity of the total potential variation of the cause variable with respect to its effect (labeled “INF,” after Woodward 2010) as well as the specificity of the actual variation in select populations, or “SAD” (in-

2. There are other attempts that all have their merits in my view, e.g., Stegmann’s (2014) idea of external ordering or Shea’s (2007) infotel theory, but they are not considered here.

spired by Waters 2007). In addition, they examine my idea of considering only the relevant potential variation (REL), where relevance is defined in terms of biological normality (Weber 2013). I take the upshot of their analysis to be as follows. First, the specificity of causes depends strongly on what probability distribution over the states of a causal variable is assumed. Second, in the context of developmental biology, SAD returns no significantly greater specificity for DNA than for alternative splicing. Third, in spite of being targeted at relevant variation, REL alone fails to provide a sufficient criterion for relevant causal specificity; a sufficient criterion would have to take into account additional parameters such as timescales.

I applaud the introduction of a precise, quantitative definition of causal specificity to this debate and agree with Griffiths et al. that our causal specificity comparisons had better be biologically meaningful. What I want to show here is that the quantitative version of causal specificity as defined by Griffiths et al. can actually do the job for which its qualitative ancestor was introduced into the philosophy of biology by Woodward and Waters: to provide a rationale for the biologists' highlighting of DNA and related biomolecules in some of their explanations.

I proceed as follows. First, I briefly present the specificity measure proposed by Griffiths et al. (sec. 2). In section 3, I examine which kind of causal specificity is meaningfully compared between different biological causes. I conclude that it is REL. Section 4 then applies this measure to alternative splicing, and section 5, to DNA sequence variation. Section 6 draws together my main points.

**2. What Is Causal Specificity, and How Can We Measure It?** Both Waters (2007) and Woodward (2010) introduce causal specificity by using Lewis's (2000) concept of influence (which the latter introduced for a different purpose, i.e., to define the causal relation itself). On the basis of this, Woodward defines what it means for a causal relation to be specific:

(INF) There are a number of different possible states of  $C(C_1 \dots C_n)$ , a number of different possible states of  $E(E_1 \dots E_m)$  and a mapping  $F$  from  $C$  to  $E$  such that for many states of  $C$  each such state has a unique image under  $F$  in  $E$  (that is,  $F$  is a function or close to it, so that the same state of  $C$  is not associated with different states of  $E$ , either on the same or different occasions), not too many different states of  $C$  are mapped onto the same state of  $E$  and most states of  $E$  are the image under  $F$  of some state of  $C$ . This mapping  $F$  should describe patterns of counterfactual dependency between states of  $C$  and states of  $E$  that support interventionist counterfactuals. Variations in the time and place of occurrence of the various states of  $E$  should similarly depend on variations in the time and place of occurrence of states of  $C$ . (Woodward 2010, 305)

For Woodward, the relation INF (after Lewis's term "influence") does not define causal dependence; it is rather a relation that causal links may or may not manifest. However, given this definition it could be argued that INF is a matter of degree where the degree depends on the number of states that are correlated by the mapping as well as the closeness of the mapping to a bijection. Griffiths et al. argue that both of these two factors affect the amount of information that is gained about the effect by intervening on a cause variable. To take this into account, they define causal specificity, or SPEC, as the information gained about the state of the effect variable by setting the cause variable to an exogenously determined value. Using information theory, Griffiths et al. identify SPEC with the difference between the entropy of the effect variable's value set  $H(E)$  and its entropy conditional on setting the cause variable to some specific value. Formally, this can be represented by using Pearl's 'do ( )' operator that is symbolized by a hat:

$$I(E; \hat{C}) = H(E) - H(E|\hat{C}).$$

The causal specificity of  $C$  relative to  $E$ , thus, is the mutual information  $I(E; \hat{C})$ , which corresponds to the difference in entropy before and after an intervention. Having briefly outlined the basic idea of the specificity measure proposed by Griffiths et al., I now turn to the different kinds of variation for which causal specificity may be compared.

**3. Which Kind of Causal Specificity Should We Compare?** As was already indicated, causal specificity can be measured over the range of all possible values that a causal variable can take (INF), some range of values that are taken to be relevant (REL), or the actual values that the variable takes in a real population (SAD). When comparing biological causes, which kind of causal specificity matters for highlighting some causes for their explanatory salience? I will briefly argue here that the most informative feature is REL.

Let us first consider SAD, the actual variation. In their comparisons involving alternative splicing, Griffiths et al. want to take into consideration the causal specificity that inheres in "the variation between cells in an organism, both spatial and temporal" (2015, 545). However, in their actual calculations they assume that the splice variants are equally probable. This seems to be at odds with their intention to use "the actual probability distribution over the values of a causal variable in some population" (541–42), for this would require that the actual frequencies of the different splice variants be taken into account. If some splice variants were overwhelmingly more abundant than others, this could significantly reduce the entropy of the probability distribution and therefore the causal specificity. It is of course unlikely that all splice variants occur with the same frequency in any actual population; thus, it is not clear what the entropy figures calculated actually mean with respect to

actual populations. Griffiths et al. cite the lack of data on the real probabilities as the main reason for this.

It seems to me that the difficulties with SAD run deeper than a lack of data. The main problem is that SAD is very sensitive to the relative abundance of a causal factor in some defined population. Thus, SAD values will be highly context dependent, to such an extent as to make any kind of systematic comparison across contexts difficult.

INF is the causal specificity of the possible range of values for a variable, which may be assumed to be equiprobable. Mad or gerrymandered possibilities show that there is a problem with this measure, which is due to a lack of constraints on the space of possibilities. For example, in Weber (2013) I discuss a scenario for protein synthesis in which the codon-amino acid assignments jointly mediated by tRNA and aminoacyl-tRNA synthase enzymes are altered after each round of the ribosome cycle by a hypothetical intervention. A similarly mad scenario exists for alternative splicing (see sec. 4). I see no reason why such mad scenarios, even though they are physically possible, should be considered as biologically meaningful because these possibilities are inaccessible. For this reason, I think the most meaningful comparisons to make are between some relevant sets of possibilities. Of course, different relevance criteria are imaginable, but my suggestion to include those possible values of the variables that could be produced by biologically normal interventions seems appropriate in at least some biological contexts (Weber 2013). Biologically normal interventions as I introduced them in Weber (2013) are such that they (1) could also be a result of *natural processes* at some nonnegligible probability and (2) are compatible with the *normal biological functioning* of the rest of the organism.<sup>3</sup> A point mutation would be an example of a biologically normal intervention in this sense, while a complete change of the codon specificities of tRNA after each round of the ribosome cycle would not be. This is why biologists consider DNA and mRNA to be information-bearing molecules while tRNAs and aminoacyl-tRNA synthase molecules are viewed as being part of the machinery that merely “reads” or “expresses” the genetic information, even though there is no difference in causal specificity of the potential variation of these biomolecules. And this is also why they consider the causal specificity of the potential DNA as well as mRNA variation to be biologically relevant, while a large range of possible splice variants, namely, those that are not producible by biologically normal interventions, are of no interest to them.

The most important alternative to this approach is Waters (2007), who has persuasively argued that biologists are often interested in actual-difference-

3. As Ken Waters pointed out to me, there are interesting similarities between my notion of biological normality and Wimsatt’s (2001) concept of generative entrenchment, to be explored elsewhere.

making causes. It should be noted that to agree to focus on actual-difference-making causes does not commit us to determining SAD in the way in which Griffiths et al. do, who measure causal specificity only over the realized states of a variable. By contrast, Waters (2007, 574) requires that only many, not all, of the different states of a variable be realized in a population. Thus, his conception of a specific actual-difference-making cause is one where a causal variable for which there is actual variation in a population and where this variation explains the actual variation in an effect variable may also have unrealized possible values that contribute to its specificity. This is different from Griffiths et al.'s SAD, and it is fully compatible with my approach taken here, so long as it is made clear that the unrealized possible values of the variable in question are in the set of relevant possibilities.

Waters's account thus construed may also be able to deal with the gerrymandered cases discussed above. For those cases require variation in causal variables that are usually not actual difference makers, for example, codon-amino acid assignments or the recognition sequences for splicing. However, Waters's account will also exclude cases such as the DSCAM gene examined by Griffiths et al. where DNA is not an actual-difference maker in the populations of interest. For this reason, I think that REL is the best choice.

A final but important desideratum, as Griffiths et al. (2015, 545) remind us, is that the relevance criteria be "rigorously enforced" for both genetic and nongenetic causes when comparing their causal specificity, on pains of violating parity of reasoning. This is what I aim for in the following two sections.

**4. Relevant Potential Variation due to Alternative Splicing.** One example considered by Griffiths et al. is the *Drosophila* DSCAM gene.<sup>4</sup> This gene has a complicated intron-exon structure and is subject to a remarkable and unusually massive amount of alternative splicing. This means that, depending on the cell's differentiation state, different parts of the gene are removed by splicing. Thus, rather than coding for a single polypeptide, the gene provides coding cassettes that can be combined in many different ways. By mutually exclusive alternative splicing,<sup>5</sup> there are an impressive 38,016 splice variants, each of which can lead to the production of a different protein molecule. While alternative splicing is quite common in eukaryotic genes, not

4. DSCAM stands for Down Syndrome Cell Adhesion Molecule. It is a surface protein expressed predominantly in neurons and appears to mediate specific cell-cell interactions. In invertebrates it is implicated in immune defense, while its role in mammals is less clear at present.

5. In this mode of alternative splicing, one of two exons is retained while the other one is removed. Other modes include exon skipping, where a number of adjacent exons are either retained or not, or alternative 5'-donor or 3'-acceptor sites.

all cases exhibit this massive range of splice variants. For the purposes of this discussion, we can describe the splice mechanism as a cause of protein sequence that has 38,016 different states that map bijectively to 38,016 different protein sequences. Using their specificity measure, Griffiths et al. calculate that the causal specificity of this mechanism amounts to 15.2 bits. At the same time, they calculate the causal specificity of the DNA to 0 bits because they assume that there is only one state in the population they are considering (a population of neurons in a *Drosophila* brain). Thus, Griffiths et al. compare the causal specificity of actual variation of the different causes (SAD) here. If this is what is being compared, alternative splicing turns out to be causally more specific than DNA variation.

How about the other examples discussed in the paper? First, let us consider the DSCAM homologs in humans. Here, there are two homologous genes, each of which has only three splice variants. Thus, the causal specificity of splicing comes out as 1.6 bits, while that of DNA sequence variation is 1 bit. Again, the causal specificity measured concerns actual variation. Finally, in the case of the entire class of vertebrate cell adhesion molecules, some 100 genes are capable of generating approximately 150 splice variants each. In this case, the causal specificity amounts to 7.2 bits for alternative splicing and 6.6 bits for DNA variation. Griffiths et al. conclude that, in this case, “both DNA and splicing variables are important determinants of diversity in this class of transcripts” (2015, 549).

What happens when we take into account not only the actual variation but also all REL? Is it true that, as Griffiths et al. (2015, 546) claim, “the machinery of splicing also changes over evolutionary time, so in the evolutionary case the ‘biologically normal’ variation in splicing is greater than the amount of variation observed in any actual population”? I must disagree. What we are considering here is the possible variation in a protein sequence producible by different alternative splicing machineries while holding constant the sequence being spliced. This includes only the variants producible by alternative splicing of this same sequence. Within these constraints, evolutionary change in the splice machinery cannot produce additional variants, or so I will argue.

To see this, we must take into account some details of the splicing mechanism (see Alberts et al. 2015, 310–20). The boundaries of exons (coding sequences) and introns (interspersed noncoding sequences) are marked by three kinds of repeated RNA sequences known as “splice signals” that are required for RNA splicing to occur: The 5'-end of introns is marked by the sequence AG|GURAGU, where R is a purine (A or G). The vertical bar indicates where the splice enzymes cut the RNA when the intron is removed. Somewhere within the intron sequence, we find the signal YURAC, where Y is a pyrimidine (C or U). Finally, at the 3'-end the signal YYYYYYYN CAG|G defines the end of the intron, again with “|” showing the exact splice site.



Only if these three signals occur in a repeated fashion on an RNA molecule can alternative splicing work. Of course, there is no necessity in the precise nucleotide sequences of these signals; they could well be otherwise. Also, there is no telling how many different such sequences could do the same job if the specificities of the splice enzymes were altered. Clearly, by altering the enzyme specificities and the splice signals, a vast number of alternative splice variants could be produced from any given gene.

Nonetheless, the question is whether this potential variation is what matters. When we consider a specific genetic locus such as *DSCAM* and want to know what the causal specificity of the alternative splicing mechanism is at this locus, then we should let the RNA sequence of the primary transcript be unaltered. Now, this sequence is very unlikely to have another set of repeated sequences that could serve as splice signals in addition to its actual splice sequences. If we altered the specificity of the splice enzymes, we would therefore not obtain a whole new set of alternative splice variants that could be made from the same gene sequence. Instead, there would not be any splicing going on at all. This is the reason why, unlike the DNA sequence and the mRNA sequence, the splice mechanism itself does not have any excess potential variation. What you see is what you get.

Having thus presented an argument that limits the variation and therefore causal specificity due to alternative splicing, I will discuss two potential sources of extra variation. First, do we not have to take into account evolutionary mechanisms such as exon repetition, shuffling, and inversion, as well as cryptic splice sites?<sup>6</sup> I think we can safely disregard the first three of these. For as I have already argued, what we are interested in here is potential variation that is due to variation in splice enzymes while we hold the primary transcript constant. The first three evolutionary mechanisms just mentioned change the primary transcript, so they are irrelevant to my argument. The case of cryptic splice sites is somewhat more complex. Here, we must distinguish between two distinct phenomena: (1) inactive splice sites that become activated due to a cis-mutation in the splice site itself (such mutations also change the primary transcript and are therefore not relevant to my argument) and (2) cases of so-called splicing error where the spliceosome erroneously recognizes a cryptic signal that resembles a normal splice site with some low frequency (see Alberts et al. 2015, 321–22). This kind of variation may be relevant to my argument; however, it is unlikely to contribute much to splicing-related causal specificity due to both its low frequency and its aleatoric character.<sup>7</sup>

6. I wish to thank an anonymous reviewer for drawing my attention to these mechanisms.

7. In general, noise in genetic mechanisms is an interesting issue from the point of view of information theory that might merit philosophical scrutiny.



A second potential source of additional potential variation that we have to consider comes from mad gerrymandered cases, such as what follows. Imagine that, each time after an intron is removed, the recognition sequences of the splicing enzymes change by a hypothetical intervention. To make this scenario more precise, consider the real causal graph that connects a cell's set of spliceosomes with its population of mRNAs. If alternative splicing is going on, the former is a causal difference maker with respect to the latter. In accordance with interventionist causal theory, we can represent the spliceosomes and the mRNAs by variables, say,  $C$  and  $E$  that take discrete values  $C_1 \dots C_n$  and  $E_1 \dots E_n$ . Now for the mad part of the scenario: instead of over well-behaved spliceosomes or different states thereof, we let the  $C$  variable range over *hypothetical* or mad splice agents that change their recognition sequence each time after cutting an intron. For example, we could let each individual value  $C_i$  stand for a mad spliceosome variant that recognizes a sequence  $\alpha$  in the first cut,  $\beta$  in the second cut,  $\gamma$  in the third cut, and so on. Now let  $C$  range over all possible combinations of  $\alpha, \beta, \gamma$  and over all other possible recognition sequences  $\delta, \epsilon, \phi \dots$  and combinations thereof. Obviously, in this way a much larger number of different mRNA sequences can be produced from the primary transcript than with well-behaved spliceosomes that recognize the same splice signals always. Because it measures potential variation and maximal entropy and the mad spliceosomes are conceptually and physically possible, INF therefore returns a high causal specificity. However, this is exactly the kind of irrelevant variation that is ruled out by the biological normality criterion. The upshot of this discussion is that INF cannot distinguish between relevant and irrelevant causal specificity.

To conclude, I have shown in this section that the REL in alternative splicing is identical to the variation producible by mutually exclusive alternative splicing and related mechanisms of the respective gene. As a result, the relevant causal specificity REL is given by the values that Griffiths et al. calculated, and we do not have to worry how many different splice mechanisms evolution might be able to produce.

**5. Relevant Potential DNA Sequence Variation.** Of course, we still need to show quantitatively that causal specificity of the REL kind is actually greater for the coding sequences. It is easy to calculate the mutual information of possible interventions on DNA or mRNA coding sequences with respect to protein sequences. Let us consider a coding sequence of 999 bases length, which is about average (DSCAM is much longer, with 6 kb). Applying combinatorics, there exist  $4^{999}$  different sequences of that length (because there are that many ways of combining the four bases A, T, G, and C to a string of 999 bases length). Because of the triplet code, this sequence can in principle code for a protein of 333 amino acids length. Because there are 20 amino acids to choose from at each position, there are  $20^{333}$  possible

protein molecules that could be made. This is much more than the number of atoms in the universe, which is estimated to be in the region of  $10^{80}$ .

Let us calculate the causal specificity of the causal connection DNA  $\rightarrow$  protein. From the  $4^{999}$  possible nucleotide sequences already calculated we can make  $20^{333}$  different proteins of that length, which equals about  $4^{720}$ . This reduction by  $4^{279}$  is due to the redundancy of the genetic code. The mRNA  $\rightarrow$  protein mapping is therefore not a bijection; it is surjective but not injective. However, for a causal specificity calculation à la Griffiths et al. this does not matter because we can assume for our purposes that the value of the mRNA sequence variable completely determines the protein sequence variable. In such a case, the mutual information about the effect variable that can be obtained by setting the cause variable to a certain value is given by the number of states of the effect variable, assuming that these are equiprobable. We thus get an information content of  $\log_2 20^{333} = 1,441.6$  bits. In other words, setting the state of the mRNA variable by an intervention reduces our uncertainty about the corresponding protein's amino acid sequence by 1,441.6 bits, which is by far superior to that of the alternative splicing mechanism.

Critics may object to these figures as not being relevant because this potential variation is surely never realized. So how is it biologically meaningful?

To counter this objection, I wish to point out that we can limit the range of variants to biologically realistic scenarios and still obtain a causal specificity that is vastly superior to anything an alternative splice mechanism could produce. An example of such a scenario would be as follows: let us consider the variation that can be produced by taking a protein of 333 amino acids length and allowing two independent amino acid substitutions at two different sites. This kind of double mutation occurs frequently enough to be biologically relevant. Importantly, it does not require an evolutionary timescale (cf. Griffiths et al. 2015, 551). Evolution does not have to enter into the picture at all. Even on much shorter timescales, for example, a few generations, the relevant protein sequence variation producible by DNA mutations is causally more specific than variation due to alternative splicing.

All that matters for my argument is that we are able to answer questions such as this: What would be a biologically relevant range of alternative states (or possible worlds, if you prefer) for this gene? Answer: All the allelic variants that could have been produced by a few biologically normal interventions at a nonnegligible probability in the immediate ancestors of the cell/organism in question, for example, point mutations. My point is that even this restricted range of alternatives quickly takes us to very high values of causal specificity.

For two point mutations, which could surely occur within a few generations in a biologically normal way, we already have a causal specificity of

$$\log_2 \left( \frac{333 \times 332}{2} \times 20^2 \right) \approx 24.4,$$

which is already substantially higher than alternative splicing in *Drosophila* DSCAM. The figure rises rapidly if we allow not only single amino acid substitutions but other naturally occurring mutations such as frameshift mutations, insertions, deletions, inversions, duplications, and so on. This kind of variation is biologically relevant. The reason why it is relevant is not the fact that it is actually realized in any population (for it may not be for all DNA sequences) but that new variants can be produced by biologically normal interventions, as it was the case for alternative splicing. I conclude that there is a relevant kind of causal specificity with respect to protein sequences that is greater for DNA and mRNA than for alternative splicing.

**6. Conclusion.** The mutual information of variables associated with interventions on a causal variable, proposed as a measure of causal specificity by Griffiths et al., distinguishes some causal factors from others and may very well be what incites biologists to often highlight DNA as a major cause, even though myriad other causal factors are involved in most biological phenomena. I have argued in this discussion note that what matters in many biological contexts is not the causal specificity of the actual variation of an actual-difference maker or the potential variation of the actual-difference-making causes but the specificity of the REL in a causal variable. And I have shown that this REL exhibits a higher causal specificity in the case of DNA and mRNA than in the case of splicing agents in biologically realistic cases.

#### REFERENCES

- Alberts, Bruce, Alexander Johnson, Julian Lewis, David Morgan, Martin Raff, Keith Roberts, and Peter Walter. 2015. *Molecular Biology of the Cell*. 6th ed. New York: Garland Science.
- Griffiths, Paul E. 2001. "Genetic Information: A Metaphor In Search of a Theory." *Philosophy of Science* 68:394–412.
- Griffiths, Paul E., and Russell D. Gray. 2005. "Three Ways to Misunderstand Developmental Systems Theory." *Biology and Philosophy* 20 (2): 417–25.
- Griffiths, Paul E., Arnaud Pocheville, Brett Calcott, Karola Stotz, Hyunju Kim, and Rob Knight. 2015. "Measuring Causal Specificity." *Philosophy of Science* 82:529–55.
- Griffiths, Paul E., and Karola Stotz. 2013. *Genetics and Philosophy: An Introduction*. Cambridge: Cambridge University Press.
- Lewis, David. 2000. "Causation as Influence." *Journal of Philosophy* 97:182–97.
- Oyama, Susan. 2000. *The Ontogeny of Information: Developmental Systems and Evolution*. 2nd rev. ed. Durham, NC: Duke University Press.
- Sarkar, Sahotra. 1996. "Biological Information: A Sceptical Look at Some Central Dogmas of Molecular Biology." In *The Philosophy and History of Molecular Biology: New Perspectives*, ed. Sahotra Sarkar, 187–231. Dordrecht: Kluwer.
- Shea, Nicholas. 2007. "Representation in the Genome, and in Other Inheritance Systems." *Biology and Philosophy* 22:313–31.
- Stegmann, Ulrich. 2014. "Causal Control and Genetic Causation." *Noûs* 48:450–64.
- Waters, C. Kenneth. 2007. "Causes That Make a Difference." *Journal of Philosophy* 104 (11): 551–79.
- Weber, Marcel. 2006. "The Central Dogma as a Thesis of Causal Specificity." *History and Philosophy of the Life Sciences* 28:595–609.

- . 2013. “Causal Selection versus Causal Parity in Biology: Relevant Counterfactuals and Biologically Normal Interventions.” In *What If? On the Meaning, Relevance and Epistemology of Counterfactual Claims and Thought Experiments*. Preprint Series 3. Konstanz: University of Konstanz.
- Wimsatt, William C. 2001. “Generative Entrenchment and the Developmental Systems Approach to Evolutionary Processes.” In *Cycles of Contingency: Developmental Systems and Evolution*, ed. S. Oyama, R. Gray, and P. Griffiths, 219–38. Cambridge, MA: MIT Press.
- Woodward, James. 2010. “Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation.” *Biology and Philosophy* 25:287–318.