

Further Clarification on Permissive and Instructive Causes

Brett Calcott

To Appear in “Biology and Philosophy”

Abstract

I respond to recent criticism of my analysis of the permissive-instructive distinction and outline problems with the alternative analysis on offer. Amongst other problems, I argue that the use of formal measures is unclear and unmotivated, that the distinction is conflated with others that are not equivalent, and that no good reasons are provided for thinking the alternative model or formal measure tracks what biologists are interested in. I also clarify my own analysis where it has been misunderstood or ignored.

1 Introduction

In the paper “Measuring Causal Specificity” (Griffiths et al. 2015), my co-authors and I outlined a formal approach to capturing the notion of causal specificity. Our proposal was to capture specificity using mutual information. Though we typically use mutual information to measure the association between two variables, when applied to variables on a causal graph it can tell us how much control we have over one variable by manipulating another. Following a suggestion of Woodward’s (Woodward 2010), I then explored whether these formal measures of specificity might help clarify the permissive-instructive distinction (Calcott 2017). Woodward’s suggestion was that instructive causes are more specific than permissive causes. I argued that a straightforward equivalence between causal specificity and the permissive-instructive distinction did not capture key features of the distinction. I then proposed an alternative way of carving up the distinction using related information theoretic measures. I captured a permissive cause with a measure of its switch-like nature that I called *precision*. And I argued that although an instructive cause was highly specific, it was not specific for an effect, but instead for controlling the relation between another cause and its effect (see the paper for details). In biological terms, fine-grained control of the DNA (the instructive cause) enables the creation of a robust developmental switch (the permissive cause) between two complex states.

In a response to this paper, Bourrat has offered his own analysis of the distinction (Bourrat 2019b). In contrast to me, he argues that the permissive-instructive distinction *is* equivalent to low-versus-high specificity. Central to Bourrat’s analysis is the following criticism. Both the “Measuring Specificity” paper and my paper on the permissive-instructive distinction treat specificity as fine-grained control, or what I shall refer to it as *range-of-influence* specificity for the rest of this paper. As Bourrat points out, there is a second concept of specificity discussed by Woodward (Woodward 2010). This is *one-to-one* specificity, where each cause has a single effect. According to Bourrat, omitting this second kind of specificity leads to a series of errors in my analysis. Once we have a way to measure one-to-one specificity, which Bourrat has recently defined (Bourrat 2019a), then the permissive-instructive distinction can be captured by combining it with a measure of range-of-influence specificity.

Bourrat is right to point out I did not address Woodward’s one-to-one measure of specificity, and that my notion of “precision” is related. However, the claim that omitting this particular measurement led my analysis astray is incorrect.

In what follows, I outline several problems in his critical paper and clarify my own position. First, I look at the measure of one-to-one specificity that Bourrat introduces, pointing out where it goes awry, and suggesting a replacement. Second, I show that his critical analysis fails to connect with the biology as he assumes the distinction is equivalent to those already discussed in philosophy. Last, I show that Bourrat mistakes the analysis of a model with the analysis of distinction, thus ignoring a central contribution of my paper.

2 Bourrat’s Many Measures, Clarified and Upgraded

To avoid a lengthy digression, in this section I assume the reader is familiar with application of information theory to measuring causation (see Griffiths et al. 2015). Griffiths and colleagues argue that range-of-influence specificity can be captured by measuring the mutual information between a cause variable and an effect variable, and call this measure *SPEC*:

$$SPEC = I(\hat{C}; E)$$

The hat on the causal variable, \hat{C} , means that we have intervened, setting this variable to a value, and then observed the subsequent change in E . This differs from merely observing C and E .

Bourrat’s account of one-to-one specificity, both in his commentary and elsewhere (Bourrat 2019a), discusses three related measures. It is important to distinguish each of them. First, Bourrat introduces Variation of Information (*VI*).

$$VI(E; \hat{C}) = H(E|\hat{C}) + H(\hat{C}|E)$$

The crucial property of *VI* (in this context) is that when zero, it shows that the mapping between cause (C) and effect (E) variables is one-to-one. Much of Bourrat’s discussion focusses on this measure. Despite this, when he compares causes in any of his models, Bourrat uses a normalised version of this measure (*NVI*):

$$NVI(E; \hat{C}) = \frac{H(E|\hat{C}) + H(\hat{C}|E)}{H(C, E)}$$

The normalised measure always lies between 0 and 1. Like VI , when NVI is zero, the relationship between the variables being measured is one to one. The shift to a normalised version is essential because, as Bourrat points out: “everything else being equal, a causal variable with a higher entropy will have a higher variation of information with the effect variable, than a causal variable with a lower entropy” (Bourrat 2019a). Normalising the measure removes this problem, allowing a comparison between variables purely on how one-to-one they are.

The switch to a normalised measure has some implications I explore below. First, however, notice that NVI works in the opposite direction to $SPEC$. When $SPEC$ is higher, range-of-influence specificity is higher, but when NVI is higher, one-to-one specificity is lower. This leads to some awkward language, with Bourrat calling it a measure of “un-specificity”, as the higher it is, the less (one-to-one) specific the relation is. Given NVI is normalised between 0 and 1, we can avoid this awkward talk by using the complement of the measure, call it $CNVI$:

$$CNVI = 1 - NVI$$

With this simple change we get all the same benefits as before, except now both range-of-influence and one-to-one specificity increase as their measure increases. When $CNVI$ is 1 (when NVI is 0), the causal relationship is one-to-one.

Using the complement has another benefit: it clarifies the relationship between range-of-influence and one-to-one specificity. Given the definition of NVI given above, we can rewrite $CNVI$ as:

$$CNVI = \frac{H(\hat{C}, E)}{H(\hat{C}, E)} - \frac{H(E|\hat{C}) + H(\hat{C}|E)}{H(\hat{C}, E)}$$

Now, that top row is equal to $I(\hat{C}; E)$, so $CNVI$ is equivalent to:

$$CNVI = \frac{I(\hat{C}; E)}{H(\hat{C}, E)}$$

We see that *CNVI* is simply a causal version of normalized mutual information.¹ Thus, the one-to-one notion of causal specificity that Bourrat actually uses (*NVI* rather than *VI*) can be substituted with normalised (causal) mutual information. Recall that *SPEC*—the measure proposed for range-of-influence specificity by Griffiths et al.—just is causal mutual information, which is the top row in the equation above, $I(\hat{C}; E)$. We now see a straightforward relationship between the two kinds of specificity: one-to-one specificity is simply a normalized version of range-of-influence specificity.

To recap, whilst Bourrat focuses his attention on the variation of information (measure 1), he uses a normalised version of variation of information (measure 2) in all of his analyses. I proposed an alternative measure (*CNVI*), which can be substituted for this second measure. This measure turns out to be a normalised version of *SPEC*. As well as having the benefit of requiring far less linguistic gymnastics to use, this alternative measure has a much clearer relationship to range-of-influence specificity.

Finally, Bourrat proposes a third measure. Both in his comment and in the paper on one-to-one specificity, Bourrat tells us that the correct way to measure causal specificity (and hence to distinguish permissive and instructive causes) is to use *both* measures. For Bourrat, high specificity (instructive) causes occur when a cause is both fine-grained (high *SPEC*) and when it is one-to-one (high *CNVI*, low *NVI*), and low specificity (permissive) causes occur when a cause has low specificity (low *SPEC*) and is not one to one (low *CNVI*, high *NVI*).

This is a perplexing move. For though both measures track something referred to as “specificity”, I see no reason to meld them together. Nor does Bourrat produce any evidence that combining them in this way captures scientists use of the term “specificity”. Furthermore, melding them in this way invites confusion, for we now have a distinction with two dimensions. It follows that there are cases where one measure may be high whilst the other is low. In fact, DNA, which is at the centre of many of these debates about causation in philosophy of science, looks like a candidate. Consider the following: there are an enormous number of different DNA sequences and a wide

¹ There is more than one way to normalize mutual information. See https://en.wikipedia.org/wiki/Mutual_information

range of downstream gene products, yet the relationship is far from one-to-one (there are synonymous mutations for a start). So it appears we have high range-of-influence specificity and low (or at least, far from maximal) one-to-one specificity. Furthermore, both of these features are important. The enormity of possibilities that DNA can support is essential for generating variation, and the lack of a one-to-one mapping between genotype and phenotype is crucial for complex evolutionary change (Wagner 2011). So the class of causal relationships that are neither high nor low specificity (at least according to Bourrat) may well be cases worthy of attention.

In contrast to Bourrat, I think it better we keep these two concepts of specificity separate, especially since we now have a clearer understanding of how to measure them, and how they relate to one another.

Before moving on, I want to discuss one last issue regarding the information theoretic measures deployed in these papers. What I call “precision” in my paper is closely related to one-to-one specificity, a point that I did not connect to Woodward’s ideas.² A causal relationship is “precise”, in my terminology, if every cause results in precisely one effect (rather than, say, one effect half the time, and a different effect the rest of the time). This isn’t necessarily one-to-one, as many causes could lead to precisely one effect. Bourrat goes on to claim that his measure is “more general” than my own, but this cannot be correct. If something is one-to-one, then it is precise, but not vice-versa. This makes the one-to-one relation more restrictive, rather than more general. Furthermore, the definition was carefully constructed to define a *robust-switch*, where many different possible states of affairs lead to just a few precise states (such as the relationship between many initial conditions and few equilibria in many dynamical systems). So although the measures are related, the notion of precision is distinct, and chosen to fit the requirements of my analysis.

² I thank Karola Stotz (pers. comm.) for pointing this out to me.

3 Nailing down the Elusive Distinction

One starting point for my paper was a comment by Woodward in his paper on causation in biology (Woodward 2010). Here, Woodward referred to the “elusive contrast” between permissive and instructive causes. Like many concepts deployed in biology (consider altruism or evolvability), nailing down precisely what the terms ‘instructive’ and ‘permissive’ mean is not straightforward. Woodward suggested that causal specificity, or something like it, was at the heart of the distinction. It looks amenable to this analysis, as two purported causes are being compared and judged differently. To do the distinction justice, however, any clarification must reflect what biologists are interested in. Thus, despite its elusive nature, we need to begin by looking at the distinction and what role it plays in biology. This should both guide our attempt to provide a definition and, once we’ve come up with a definition, verify whether it delivers the right results. So how does Bourrat understand the distinction and how does he connect his formal work to the distinction?

Bourrat begins by equating the permissive-instructive distinction with a distinction from philosophy: the notion of background and triggering conditions. A commonly used example that captures this distinction is lighting a match: the striking of the match is the triggering condition and oxygen is the background condition. In addition, he tells us that a background condition is typically less specific and triggering conditions more specific. Bourrat hedges these equivalences somewhat (making it less clear what the point of making them is) but the general picture he provides is this:

permissive = background = less specific

instructive = triggering = more specific

Both equivalences confuse the picture, rather than clarify what biologists are interested in. First, consider the idea that background causes are permissive and triggering causes are instructive. The permissive-instructive distinction originated around understanding *induction*, when the introduction of some substance results in a particular developmental pathway being taken. Induction looks like a clear case of a triggering cause, occurring against some set of background developmental conditions. What biologists have been interested in understanding is whether this

triggering cause (the inductor) is permissive *or* instructive. If we look at the history of embryonic induction, we see that, initially, biologists thought that triggers for particular kinds of complex organisation in the embryo were instructive. In the end, however, they turned out to be permissive (see the Kirschner and Gerhart quote below). The two distinctions are thus not the same. This is clear even in the quotations provided by Bourrat:

Such findings suggest that some substrate molecules may be permissive, inducing neurite elongation, but not instructive, in the sense of directing growth in a particular direction along a pathway or indicating which way to go at an intersection between two pathways (Lemmon et al. 1992).

Here, the author judges the substrate molecules as either permissive or instructive depending on whether they merely permit neurite growth or instead direct that growth in a particular way. If Bourrat's equivalence held, then we would have the odd consequence that these molecules become a background or triggering cause depending on whether we judge them permissive or instructive. But this is out of step with how biologists think about induction; these molecules are a triggering cause, regardless of what specificity we end up assigning them.

What about the idea that background causes are less specific and triggering causes are more specific? Here, Bourrat begs the question. Recall that a central aim of my paper and Bourrat's response was to question whether the instructive / permissive distinction lined up with high / low causal specificity. If permissive causes really are background causes *and* background causes really are less specific, then Bourrat has neatly defined things so that his view is already true.

Bourrat's presentation of the distinction has not fared well. How does he show his measures match the use of the distinction?

Bourrat provides several examples of biologists deploying the distinction. He then ties these usages to his formal definition in the following passage:

A permissive cause has both a lower causal range of influence (ideally two possible states, to act as a switch) and a lower one-to-one specificity (higher normalized

variation of causal information) than the other causal variables producing the same effect. These other causal variables are instructive relatively to the first cause for they have a high(er) mutual causal information which shows they are not a switch and a high(er) one-to-one specificity, which shows that they direct a particular outcome. These properties match the definitions of permissive and instructive cause used by biologists ...

The claim that his analysis “matches the definitions” of the biologists requires a leap of faith, however. To begin with, the statements quoted are not definitions, but usages of the terminology. Furthermore, the quotations do not contain terms such as “range of influences” or “causal variables” or the like, so there is no obvious way to translate their talk to any formal model of causation. To show that his analysis captures what biologists are interested in, there has to be some interpretation of these passages that links the particular examples to the proposed formal definitions. Yet what follows this statement is further discussion of how the formal measures apply to the model rather than any attempt to connect these measures to what biologists have said. Nor would establishing this connection be an easy task. For part of the elusive nature of the distinction is that biologists have used the term loosely. It is not clear at all that the selected quotes pick out a single idea.

Given this lack of clarity, how do I just justify my interpretation of the distinction? Rather than attempting to interpret a variety of usages, I focused on the origins and history of the distinction as it was applied to embryonic induction. Here, we see that the crucial features of the distinction—those that distinguish it from other causal distinctions—are intertwined with explaining how a simple switch can generate a complex outcome. Here is a concise summary from Kirshner and Gerhart:

... it seemed unlikely to many biologists that the signals of embryonic induction could act permissively. The outcome of embryonic induction is complicated; it results in creation of virtually the entire organized embryo, with hundreds of cell types and many organs, including the entire nervous system, all in the right places. The name

“organizer” that Spemann gave to the source of the inducer implied instruction, maybe even micromanagement. To everyone’s surprise, embryonic induction turned out to be a permissive process; the organizer provides a signal of little complexity (Kirschner and Gerhart 2006)

This passage highlights three key components of the distinction. First, we observe a reliable cause (induction) of some complex effect (embryonic organisation). Second, the ability of this cause to switch on this complex effect leads us to believe that the cause itself contains the means to produce the complexity of the effect (that is, the cause appears to be instructive). Third, somewhat surprisingly, the cause turns out to be something simple; it is a mere switch of turning on existing complexity (that is, the cause is actually permissive). These three lead to a fourth component: if this reliable cause is just a switch, then the cause of this complexity must lie elsewhere. Typically, elsewhere means *in the DNA* (as Scott Gilbert suggests in the citation in my paper). My analysis, including the model I constructed and the use of the formal tools, was directed to capturing these four components; to showing how a simple switch-like cause can lead to such a complex effect, and how we can make sense of the complexity lying elsewhere.

Highlighting these four components of the permissive-instructive distinction also links it to another, more central, debate in philosophy of biology concerning the idea that information is “in” the DNA. This connection is important for, as I state in the beginning of the paper, my aim was not just to apply formal causal tools to clarify the permissive-instructive distinction. A second aim was to show how we can extend our thinking about specificity in DNA beyond just coding for proteins to encompass the conditional expression of these proteins. This is a novel contribution to an old debate, for it suggests that fine-grained control over conditional expression is one way to make sense of DNA encoding something like a program (see the “Causal specificity and gene regulation” section in the paper for more on this).

4 On Models and Mischaracterisation

In the previous section I outlined some challenges faced when attempting to clarify concepts in biology. Formal methods provide one useful tool for addressing these challenges, as they provide

a concise language to express and test commitments and relationships that are not obvious in ordinary language. In one elegant example of this approach, Ben Kerr and his co-authors formalise altruism in several ways, and this explains some controversy in the cooperation literature (Kerr, Godfrey-Smith, and Feldman 2004). Both Bourrat and I make use of formal measures of causation in our attempts to clarify the permissive-instructive distinction. Though I am confident this approach can bring some much-needed clarity to the distinction, we should be clear as to role these formal measures play. For while it is possible to apply such measures directly to empirical work, here we are applying them to simple models. The aim is to construct a simple model which generates, as best we can, the same intuitions as the biological cases. Then, we apply the formal measures to the model, fine-tuning them so they capture the biological concepts in play. The role of the model and its analysis is thus as a thinking tool, rather than an empirical example, and any progress in clarifying the concepts relies on the following set of assumptions:

1. **If** the model *plausibly* represents the relevant biology, and
2. **If** the formal measure tracks something in the model that *plausibly* reflects the concept biologists are interested in,
3. **Then**, we gain some insight into how to characterise what biologists are interested in.

Step 1 is crucial here. For no matter how insightful the analysis of a model is, if that model does not represent the relevant biology, we have not made progress on the problem (Levy 2011). This point is relevant here for I begin my paper by analysing a simple model I call “Woodward’s radio”, comprising an on-off switch and a tuning dial which together control the output of the radio. I then state that, though this model may be apt for clarifying the notion of causal specificity, it is insufficiently rich to capture the permissive-instructive distinction. That is why I spent much of my paper developing a new model, one capable of capturing the components I outlined in the previous section. Bourrat, however, thinks further analysis of this same simple model is all that is called for. His justification is that:

[...] Woodward's radio is formally equivalent to the bimolecular example of transcription presented in Griffiths et al. (2015). As such, Woodward's radio also has some biological relevance.

Some biological relevance is far too weak. The model must be relevant to the issue at hand, namely defining the permissive-instructive distinction, and there is no attempt to counter my claims of its inadequacy. In addition, Bourrat's attempt to justify the model's relevance by appealing to its prior use in a paper on causal specificity also appears to presume that the permissive-instructive distinction is equivalent to causal specificity. But (again) this is the very point in question.

To see why the radio model is inadequate, we need to think about how the model connects to biology. Recall that each causal variable in the model is meant to capture something that can be manipulated. One causal variable must capture induction—that some substance is present or not, and its presence or absence switches between two distinct effects. What about the other variable? The assumption I have made is that this other variable is “in the genes”—in this case what varies is the DNA (see the discussion on “two dimensions of variation” in my paper). At a first pass, the radio looks like it can capture this—the on-off switch is like the inductor, and the tuning dial is like the genes. This is an inviting simplification, but it fails to represent the interaction between induction (manipulating one variable) and changes in the DNA (manipulating the other variable). In a radio, the on/off switch is reliable (or “precise”, as I define it in the paper). When it is off, the radio plays no sound; when it is on, the radio plays whatever channel we choose. An unreliable switch, in contrast, might only sometimes turn the radio on (say eighty percent of the time). We take the reliability of the switch for granted in a radio and hence also in the model. But in biology, the ability to switch between two very different complex states of affairs using a single causal factor is something in need of explanation. We need to explain not just the complexity of resultant embryonic organisation, but also why it is possible to so easily switch between two versions of this organisation by wiggling a single causal variable. After all, it was the capacity to switch so easily between the two states that led biologists to (incorrectly) think the inductive cause itself contained the requisite instructions for building complexity. The problem is that, in Woodward's radio model, the switch is always a switch. In contrast, in the model I introduced (the Waddington

Box), this switch-like capacity is under the control of the other variable. This ability for manipulations to the DNA to change how what a switch does and how reliable it is also makes sense biologically. For changes to the DNA can affect not only what is expressed, but also the conditions under which they are expressed when the DNA plays a regulatory role (such as a cis-binding region).

Bourrat does attempt to address the simplicity of Woodward's radio in a lengthy footnote. He suggests a simple extension to Woodward's radio model make it capable of modifying the robustness of the switch. This move just adds to the problem, as it introduces a third causal variable—one that directly controls how reliable the switch is. It is unclear how this would be dealt with, nor what biological significance it has. Even with this addition, however, the radio model is incapable of capturing essential features present in the Waddington Box model. To clarify, here is what a radio would have to be like for it to have comparative features. Such a radio might play classical music when the switch is ON, and jazz when the switch is OFF. Twiddling with the tuning dial might then make the radio play jazz half the time and classical the other half when the switch is ON, and play BBC news when the switch is OFF. Twiddle the tuning dial again, and it would play BBC news ten percent of the time when switched OFF, jazz the other ninety percent, and classic rock when switched ON. In general, for the radio to have an equivalent capacity, the tuning the dial must both a) make the switch more or less reliable, and b) cause the switch to choose between a variety of different channels.

I have called this a radio, but the usefulness of the metaphor is lost at this point—no radio is like this. The Waddington box model, in contrast, makes this all plain. I begin with a well-known physical model for explaining probability distributions and connect it to influential ideas about developmental stability and evolution; the model visually demonstrates causal pathways via the route a ball takes through the box; and, though significantly more complex than a radio, the model is amenable to a complete analysis using information theoretic measures.³ In addition, the model

³ It is interesting to note what Bourrat's own method of analysis produces when applied to this model. As I show in the paper, range-of-influence specificity is low for both variables in a simple competitive analysis. Thus, we should conclude, following Bourrat, that neither cause is instructive. Yet, as I show, wiggling the pins produces over three million different layouts and over five thousand different ways of distributing the balls. Only a hierarchical analysis, as I propose, captures this.

demonstrates why the flexibility over the switching mechanism (which is absent in the radio model) produces a kind of specificity relevant to evolutionary change. For a crucial feature of a regulatory system is that it can wire up existing cues to many different complex outputs in just the right way. This point is so important that it is one of the key aspects of Kirschner and Gerhart's theory of facilitated variation: "weak linkage" (Kirschner and Gerhart 2006).

5 Bourrat's Hammer

I've outlined numerous gaps and problems in Bourrat's analysis, but perhaps the core difference between our approaches is what we take to be elusive about the permissive-instructive distinction. For me, what is elusive is how to best make sense of the distinction given its origins and usage in biology. In contrast, for Bourrat, what eludes us is merely how to measure it. He presumes that the distinction itself is clear, as he sees no problem in citing several usages of the distinction without further clarification. He also presumes that the distinction is adequately reflected in the simple radio model, for he deploys this model with little reflection on its limitations. Thus, for Bourrat, what remains is merely to figure out the right way to measure the causal interactions in the model. Fortunately, he has just the measure at hand: his own newly minted measure of one-to-one specificity (Bourrat 2019a). As I have shown, this approach forces the biology into an existing causal distinction rather than furthering our understanding of the permissive-instructive distinction.

6 Acknowledgements

The initial ideas for much of this work occurred during my time in the Theory and Method in Biosciences group at the University of Sydney with Paul Griffiths, Arnaud Pocheville, and Karola Stotz, where funding was provided by Swansea University Templeton World Charity Foundation. Many thanks to Maureen O'Malley and Emily Parke for reading and commenting on drafts.

7 Bibliography

- Bourrat, Pierrick. 2019a. "Variation of Information as a Measure of One-to-One Causal Specificity." *European Journal for Philosophy of Science* 9 (1). <https://doi.org/10.1007/s13194-018-0224-6>.
- . 2019b. "On Calcott's Permissive and Instructive Cause Distinction." *Biology & Philosophy* 34 (1). <https://doi.org/10.1007/s10539-018-9654-y>.

- Calcott, Brett. 2017. "Causal Specificity and the Instructive–Permissive Distinction." *Biology & Philosophy*, May, 1–25. <https://doi.org/10.1007/s10539-017-9568-0>.
- Griffiths, Paul E., Arnaud Pocheville, Brett Calcott, Karola Stotz, Hyunju Kim, and Rob Knight. 2015. "Measuring Causal Specificity." *Philosophy of Science* 82 (4): 529–55. <https://doi.org/10.1086/682914>.
- Kerr, Benjamin, Peter Godfrey-Smith, and Marcus W. Feldman. 2004. "What Is Altruism?" *Trends in Ecology & Evolution* 19 (3): 135–40. <https://doi.org/10.1016/j.tree.2003.10.004>.
- Kirschner, Marc W., and John C. Gerhart. 2006. *The Plausibility of Life: Resolving Darwin's Dilemma*. New Haven: Yale University Press.
- Lemmon, V., S. M. Burden, H. R. Payne, G. J. Elmslie, and M. L. Hlavin. 1992. "Neurite Growth on Different Substrates: Permissive versus Instructive Influences and the Role of Adhesive Strength." *Journal of Neuroscience* 12 (3): 818–26. <https://doi.org/10.1523/JNEUROSCI.12-03-00818.1992>.
- Levy, Arnon. 2011. "Game Theory, Indirect Modeling, and the Origin of Morality." *The Journal of Philosophy* 108 (4): 171–87.
- Wagner, Andreas. 2011. *The Origins of Evolutionary Innovations: A Theory of Transformative Change in Living Systems*. 1 edition. Oxford ; New York: Oxford University Press.
- Woodward, James. 2010. "Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation." *Biology & Philosophy* 25 (3): 287–318. <https://doi.org/10.1007/s10539-010-9200-z>.