

When to expect violations of causal faithfulness and why it matters

Holly Andersen

Simon Fraser University

holly_andersen@sfu.ca

Abstract: I present three reasons why philosophers of science should be more concerned about violations of causal faithfulness (CF). In complex evolved systems, mechanisms for maintaining various equilibrium states are highly likely to violate CF. Even when such systems do not precisely violate CF, they may nevertheless generate precisely the same problems for inferring causal structure from probabilistic relationships in data as do genuine CF-violations. Thus, potential CF-violations are particularly germane to experimental science when we rely on probabilistic information to uncover the DAG, rather than already knowing the DAG from which we could predict the right experiments to ‘catch out’ the hidden causal relationships.

Wordcount, including references, abstract, and footnotes: 4973

1. Introduction

Several conditions must be met in order to apply contemporary causal modeling techniques to extract information about causal structure from probabilistic relationships in data. While there are slightly different ways of formalizing these requirements, three of the most important ones are the causal Markov, causal modularity, and causal faithfulness conditions. Potential failures of the first two of these conditions have already been the subject of discussion in philosophy of science (Cartwright 1999, 2002, 2006; Hausman and Woodward 1999, 2004; Steel 2006; Mitchell 2008; Woodward 2003, 2010). I will address failures in the third condition, causal faithfulness, and argue that failures of this condition are likely to occur in certain kinds of systems, especially those studied in biology, and are the most likely to cause trouble in experimental settings.

Faithfulness is the assumption that there are no precisely counterbalanced causal relationships in the system that would result in a probabilistic independence between two variables that are actually causally connected. While faithfulness failures have been discussed primarily in the formal epistemology literature, I will argue that violations of faithfulness can impact experimental techniques, inferential license, and issues concerning scientific practice that are not exhausted by the formal epistemology literature.

In particular, a formal methodological perspective might suggest a distinction between genuine and merely apparent failures of CF, such that supposed examples of CF-violating systems are not ‘really’ CF-violating, but merely close. But as I will argue, this

distinction is not epistemically justifiable in experimental settings: we cannot distinguish between genuine and merely apparent CF violations unless we already know the underlying causal structure; without this information, merely apparent and genuine CF violations will be indistinguishable. Violations of CF faithfulness are particularly germane to experimental science, since CF is the assumption that takes us from probabilistic relationships among variables in the data to the underlying causal structure. In contrast, for instance, the Causal Markov condition takes us from causal structure to predicted probabilistic relationships. Going from data to underlying causal structure is the most common direction of inference from the epistemic vantage point of science. Rather than beginning by knowing the true causal graph of the system in question to predict probability distributions, experiment moves from probabilistic relationships to the underlying causal structure.

This means that failures of CF arguably have the most potential for wreaking havoc in experimental settings, and have interesting methodological consequences for the practice of science: we should expect to find epistemic practices that compensate for CF-violations in fields that study systems where faithfulness is likely to fail. Thus, these conditions are of interest not only to those working on formal modeling techniques, but also to broader discussions in philosophy of science, especially those that concern epistemic practices in the biological, cognitive, or medical sciences.

2. Violations of the Causal Faithfulness Condition

Violation of CF occurs when a system involves precisely counterbalanced causal relationships. These causal relationships appear “invisible” when information about

conditional and unconditional probabilities is used to ascertain a set of possible causal directed acyclic graphs (DAGs) that are consistent with data from that system. More precisely:

Let G be a causal graph and P a probability distribution generated by G . $\langle G, P \rangle$ satisfies the Faithfulness Condition if and only if every conditional independence relation true in P is entailed by the Causal Markov Condition applied to G . (Spirtes, Glymour, and Scheines 2000, 31)


One can think of faithfulness as the converse of the Causal Markov condition: faithfulness says that given a graph and associated probability distribution, the only independence relations are those that follow from the Causal Markov condition alone and not from special parameter values... (Woodward 2003, 65)

Informally, variables should only be probabilistically independent if they are causally independent in the true causal graph; when causal relationships cancel each other out by having precisely counterbalanced parameter values, the variables are probabilistically independent, but not causally independent. Thus, in systems that have CF-violating causal relationships, the probabilistic relationships between variables include independencies that do not reflect the actual causal relationships between those variables.

Probabilistic relationships are used to generate possible causal graphs for the system. There may be multiple distinct causal graphs which all imply the observed set of

probabilistic relationships. The candidate graphs can then be used to generate further interventions in the system that will distinguish between the graphs; if two candidate graphs make different predictions for the consequences of an intervention on variable A, then performing this intervention on A should return an answer as to which of the candidates graphs matches the observed results. The use of probabilistic data to generate candidate causal graphs that can then be used to suggest further interventions can save huge amounts of time and energy by focusing on a few likely candidates from an indefinitely large number of candidate causal structures.

DAGs of causal faithfulness violations may take several forms. For example:



Some authors (Pearl 2000, Woodward 2010) rely on a stronger constraint, causal stability, which requires that probabilistic independence relationships be stable under perturbation of parameter values across some range, to eliminate “pathological” (i.e. CF-violating) parameter values.

Definition 2.4.1 Stability:

Let $I(P)$ denote the set of all conditional independence relationships embodied in P . A causal model $M = \langle D, \Theta \rangle$ generates a stable distribution if and only if $P(\langle D, \Theta \rangle)$ contains no extraneous independences – that is, if and only if $I(P(\langle D, \Theta \rangle)) \subseteq I(P(\langle D, \Theta' \rangle))$ for any set of parameters Θ' . (Pearl 2000)

Violating causal stability would require a system to respond to changes in one parameter value with compensating changes in another parameter, so that the values remain exactly counterbalanced for some range of values.

The potential for CF-violations to reduce the reliability of methods for extracting causal structure from data is well-known in formal epistemology. However, I will argue that philosophers of science in general should pay more attention to such violations; understanding the difficulties that CF-violations pose will enhance our ability to accurately characterize features of experimental practice, and should be included in normative considerations regarding evidence and inference. The main arguments in this paper can be summarized in three brief points:

(1) Even if CF-violating systems are measure 0 with respect to the set of causal systems with randomly distributed parameter values, this does not imply that we will only encounter them with vanishing probability. CF-violating systems may be of particular interest for modeling purposes compared to non-CF-violating systems, in particular because certain kinds of systems may have structural features that render CF-violating parameter values more likely.

(2) As an example of point 1, structural considerations regarding dynamically stable systems that are the result of evolutionary processes should lead us to expect CF-violations in various biological systems. For systems that have evolved to maintain stable equilibrium states against external perturbation, we should also expect violations of the stronger condition, causal stability. I briefly present an example of this: mechanisms for salinity resistance in estuary nudibranchs.

(3) ‘Apparent’ CF-violations in equilibrium-maintaining systems can be generated in certain experimental conditions even though the actual causal relationships in question may not be exactly balanced. Some measurement circumstances will result in a data set that violates CF, even if the actual system being measured does not genuinely violate CF. We should be as concerned with merely apparent as with genuine CF-violations, since both kinds of violations lead to the same difficulties for moving from probabilistic relationships in data to accurate DAGs of systems.

These three points highlight why philosophers of science in general should be concerned: causal systems may not genuinely violate CF, but yet pose the same problems for experimental investigations as if they did. Apparent CF-violations occur when systems do not in principle violate CF but appear to due to measurement issues connected with data-gathering. In both genuine and merely apparent CF-violations, probabilistic relationships in the data will suggest a set of candidate causal graphs that are inaccurate; as a result, further interventions will yield conflicting answers. Scientists could in principle ‘catch out’ these merely apparent CF-violations *if they knew exactly how to test for them*. But to

do this, they would need the DAG, and this is the information that they lack when proceeding from the data to underlying causal structure. When we have incomplete knowledge of the causal structure of the system under investigation, we lack this ability to distinguish between merely apparent and genuine CF-violations. Both raise the same problems.

3. The measure of CF-violating systems

Spirtes, Glymour, and Scheines (2000) offer a proof that CF-violating systems are Lebesgue measure 0 with respect to possible causal systems, while non-CF-violating systems are measure 1. “The parameter values—values of the linear coefficients and exogenous variances of a structure—form a real space, and the set of points in this space that create vanishing partial correlations not implied by the Markov condition have Lebesgue measure 0” (41). From this, they conclude that we are vanishingly unlikely to encounter CF-violating systems, and so proceed on the initial presumption that any given causal system is not CF-violating. This proof may be part of the reason why comparatively little attention has been paid to causal faithfulness compared to the causal Markov and modularity conditions. However, the fact that CF-violating systems are measure 0 in this class does not imply that we will not encounter them with any frequency.

To motivate this, consider an analogy with rational numbers. They are also measure 0 with respect to the real numbers, while irrational numbers are measure 1. And, there are circumstances under which we are vanishingly unlikely to find them. If a random real number were to be chosen from the number line, the probability that we will

draw an irrational number is so overwhelming as to warrant ignoring the presence of rational numbers. However, this does not imply that rational numbers are unlikely to be encountered *simpliciter*: bluntly put, we don't 'encounter' the numbers by randomly drawing them from the number line. Rational numbers are encountered overwhelmingly more often than one would expect from considering only the proof that they are measure 0 with respect to real numbers.

The Spirtes, Glymour, and Scheines proof assumes that all parameter values within the range of a continuous variable are equally probable (Zhang and Spirtes 2008). Without this assumption, one can't presume that the CF-violating values are vanishingly unlikely. For instance, this assumption does not hold for systems that involve equilibrium-maintaining causal mechanisms. Such mechanisms work to maintain counterbalanced parameter values, rendering it much more likely that parameter values will result in CF-violations.

It is true that if causal systems took on parameter values randomly from their range, we would expect to encounter CF-violating systems with vanishingly small probability, and in that scenario, we could safely ignore CF-violations as a real possibility on any given occasion. However, some systems survive, and become scientifically interesting targets for investigation, precisely because they achieve long-term dynamic equilibrium using mechanisms that rely on balanced parameter values. In such systems, the parameter values are most certainly not indifferently probable over their range. In fields like biology, neuroscience, medicine, etc., we are disproportionately interested in modeling systems that involve equilibrium maintaining mechanisms. This suggests that our modeling interests are focused on CF-violating systems in a way that is

disproportionate to their measure when considered against all possible causal systems. Thus, we cannot conclude from the fact that CF-violating parameter values have measure 0 with respect to all possible parameter values that we will not encounter such violations on a regular basis.

Zhang and Spirtes (2008) discuss some circumstances in which systems may violate CF. However, their discussion makes it seem like CF-violations occur primarily in artificial or constructed circumstances. One such example is homeostatic systems, which maintain equilibrium against some range of perturbations, such as thermostats maintaining a constant temperature in a room. Zhang and Spirtes demonstrate that CF can be replaced with two distinct subconditions, that, taken together, provide almost the same inferential power as causal faithfulness. If systems violate only one of these subconditions, such violations can be empirically detected. This is an extremely useful result, and increases the power of Bayes' nets modeling to recover DAGs from data. However, this result should not be taken as resolving the problem.

In particular, their use of a thermostat as example of a homeostatic system does not do justice to the incredibly complex mechanisms for homeostasis that can be found in various biological systems. Considering these more sophisticated examples provides a clearer view of the potential problems involved in modeling such systems under the assumption of causal faithfulness.

4. Evolved dynamical systems and equilibrium-maintaining mechanisms

The tendency for evolved systems like populations, individual organisms, ecosystems, and the brain to involve precisely balanced causal relationships can be easily

explained by the role these balanced relationships play in maintaining various equilibrium states (see, for instance, Mitchell 2003, 2008). Furthermore, the mechanisms by which organisms maintain internal equilibrium with respect to a huge variety of states will need to be flexible. They need to not simply maintain a static equilibrium, but respond to perturbation from the outside by maintaining that equilibrium. This means that many mechanisms for equilibrium maintenance will have evolved to keep an internal state fixed over some range of values in other variables, not merely for a single precise set of values. Any system that survives because of its capacity to maintain stability in the face of changing causal parameters or variable values will be disproportionately likely to display CF-violating causal relationships, and, more strongly also violate causal stability.

An intriguing example is nudibranchs, commonly known as sea slugs (see especially Berger and Kharazova 1997). Many nudibranchs live in ecosystems such as reefs, where salinity levels in the water change very little. These nudibranchs are stenohaline: able to survive within a narrow range of salinity changes only. In cases where salinity levels vary over narrow ranges, nudibranchs respond to changes in salinity levels by a cellular mechanism for osmoregulation, where cells excrete sodium ions or take in water through changes in cell ion content and volume. This mechanism provides tolerance, but not resistance, to salinity changes, because it maintains equilibrium by exchanging ions and water with the surrounding environment. In cases of extremely high or low salinity, this mechanism will cause the animal to extrude too much or take in too much (this is why terrestrial slugs die when sprinkled with salt).

Euryhaline nudibranchs, found in estuary environments where saline levels may vary dramatically between tides and over the course of a season or year, display a much

higher level of resistance to salinity changes. There is a pay-off, in the form of increased food sources with reduced competition for nudibranchs that are able to withstand the changing saline levels. But in these environments, the osmoregulatory mechanism for salinity tolerance is insufficient. A further mechanism has evolved in nudibranchs (and in molluscs more generally) for salinity resistance in conditions of extreme salinity variations in the external environment. These two mechanisms for salinity regulation in euryhaline nudibranchs are fairly independent. The osmoregulation mechanism is supplemented with an additional mechanism which involves hermeticization of the mantle, which prevents water and ion exchange with the outside environment.. This can accommodate changes in salinity that take place over fairly short periods of time, since salinity levels can change dramatically over the course of an hour. Instead of maintaining blood salinity at the same level as the outside environment, this additional mechanism allows the organism to maintain an internal salinity level that differs from that of its environment. Mantle hermeticization and osmoregulation are distinct mechanisms, but in contexts of extremely high or low salinity, they will both act such that the variables of external and internal salinity are independent

Further, there are two distinct mechanisms in muscle cells that work in coordination in extreme salinity cases to maintain a balance of ions inside the muscle cell. The concentration of these ions, especially sodium and potassium, can change dramatically in low or high salinity levels. There are two ion pumps in the cell that maintain overall ion concentration at equilibrium across a fairly substantial range of salinity variation in the external environment. Even though external salinity has several causal effects on the internal ion balance of a cell, these two variables will be probabilistically independent for

a range of external salinity values (in particular, for the range in which the organisms are naturally found).

The ion balance of muscle cells during adaptation to various salinities could not be achieved by virtue of the Na/K-pump alone, removing sodium and accumulating potassium. As it is clear from the data obtained, the concentration of both ions drops at low salinity and increases at high salinity. Therefore, the effective ion regulation in molluscan cells can be provided only by cooperative action of two pumps – the Na/K-pump and Na,Cl-pump, independent of potassium transport. (Berger and Karazova 1997, 123-4)

There are several points that this example illustrates. The first is that of the comparative probability that a complex system, such as an organism like a nudibranch, will display CF-violating causal relationships in the form of mechanisms that maintain equilibrium. Consider the (Spirtes, Glymour, and Scheines 2000) proof that assumes that all parameter values are equally likely. We can see how this falls apart in the case of evolved systems. Let's grant that, in some imaginary past history, all the parameter values for mechanisms such as these two ion pumps were equally likely. This would have resulted in a vast number of organisms that ended up very rapidly with internal ion imbalances and then (probably rather immediately) died. The organisms that managed to stick around long enough to leave offspring were, disproportionately, those with mechanisms that were precisely counterbalanced to maintain this internal equilibrium. Having CF-violating mechanisms would be a distinct advantage. The same applies for

other important equilibrium states –organisms with less closely matched values are less capable of maintaining that equilibrium state. Insofar as these are important states to maintain, it becomes extremely probable that. Over time, those with the closest matches for parameter values will be more likely to survive. Thus, even if we grant the assumption (already unlikely in this context) that all parameter values start out as equally likely, we can see how rapidly the CF-violating ones would come to be vastly overrepresented in the population.

The second point it illustrates is how such sophisticated equilibrium-maintaining mechanisms can violate CF in a much more problematic way than the comparatively simplistic thermostat example considered by Zhang and Spirtes.¹ Finally, note that the two ion pump mechanisms are not balanced merely for a single external salinity value: they are balanced for a range of values. Thus, this example violates not merely CF but also the stronger condition of causal stability.²

I am certainly not claiming that all causal relationships in such systems will violate CF or causal stability. But it is possible that, for any given system that involves equilibrium-maintaining mechanisms, and especially for those with sophisticated evolved equilibrium-maintaining mechanisms, there will be at least some causal relationships in

¹ Note that a DAG representing the two mechanisms for the ion pumps, connecting external salinity levels as a variable to a variable representing internal ion balance in muscle cells, is not of the triangular form that is potentially detectable using the methods in Zhang and Spirtes (2008).

² This example also provides weight to the Russo-Williamson thesis, that information about probabilistic relationships requires supplementation with information about underlying mechanisms in order to justify causal claims. These examples suggest how investigation into mechanisms for equilibrium-maintenance compensate for the methodological issues that CF violations generate; we would expect the Russo-Williamson thesis to hold particularly of systems liable to violate CF.

the system that violate either or both of these conditions. This changes the stance we take at the beginning of an investigation: rather than starting from the assumption that CF-violations are vanishingly unlikely, and only revisiting this assumption in the face of difficulties, we should start investigations of such systems with the assumption that it is highly likely that there will be at least one such spurious probabilistic independence.

5. Apparent CF-violations and their experimental consequences

Consider a possible response to the argument in the previous section. One might be concerned that the examples I offer do not involve genuine CF-violations—when examined more closely, it may turn out that the causal relationships in questions are not exactly balanced, but merely close. This response might involve the claim that even in the case of biological systems, CF is not genuinely violated, because there are slight differences in parameter values that could be identified, especially if one performed the right interventions on the systems to ‘catch out’ the slight mismatch in parameter values. Or, by taking recourse to causal stability, one might say that while the equilibrium state of some systems involves precisely counterbalanced causal relationships, in the case of perturbation to that equilibrium, these relationships will be revealed. Perturbation of systems that return to equilibrium would thus be a strategy for eliminating many (or most) merely apparent CF-violations.

Answering this challenge brings us to the heart of why CF-violations deserve broader discussion. Considered from a formal perspective, there is a deep and important difference between systems that actually violate CF, or causal stability, and those that do not. This fact motivates a response to merely apparent CF-violations that takes them to be

not methodologically problematic in the same way that genuine ones are. But the ways in which merely apparent CF-violations can be ‘caught out’ generally will require information about the DAG for the system, in order to predict precisely which variables should be intervened on, within what parameter ranges, in order to uncover closely-but-not-exactly matched parameter values. While it is in principle possible to do this, it requires knowing precisely which intervention to perform, and it is this information that will be lacking in a large number of experimental situations where we don’t already have the DAG for the system, since that is what we are trying to find.

Thus, a particular data set drawn from a target system for which investigators are seeking the DAG may have spurious conditional independencies between variables (i.e. violate CF) even though in the true DAG, those parameters are not precisely balanced. In other words, depending on how the data is obtained from the system, the data set may violate CF even though the system itself doesn’t. How could this happen? There are a soberingly large number of ways in which a data set can be generated such that a merely apparent CF-violation occurs. The point to note here is that *merely apparent violations will cause exactly the same problems for researchers as would genuine CF-violations*. There are methodological issues in dynamically complex systems such that a non-CF-violating system may nevertheless result in a dataset that is CF-violating. Here are some ways in which this may happen.

The first is quite obvious: parameter values that are not exactly opposite may nevertheless be close enough that their true values differ by less than the margin of error on the measurements. Consider the parameter values in diagram 1a. A genuine CF-violation will occur if $a = bc$. However, an apparent CF-violation will occur if $a \pm \epsilon_1 =$

$bc \pm \epsilon_2$. Concerns about the precision of measurements and error ranges are well-known, but it is useful to consider them here with respect to the issue of causal faithfulness as another way to flesh out their role in investigatory practices.

Two other ways in which apparent CF-violations may occur concern temporal factors which may play a key role in the ‘catching’ of equilibrium-balanced causal relationships. Temporal factors can distinguish systems with or without causal stability, for instance, a CF-violating system that is fragiley balanced.

Consider the time scale of a system that involves balanced causal relationships for the purposes of restoring and maintaining some equilibrium state: this may be on the order of milliseconds for some cellular processes, tens to hundreds of milliseconds for many neurological processes, minutes to days for individual organisms. After a perturbation takes place, the system will re-establish equilibrium during that range of time. In order to successfully ‘catch’ the counterbalanced causal relationships in the act of re-equilibrating, the time scale of the measurements must be on a similar or shorter time scale. If the time scale of measurements is long with respect to the time scale for re-establishing equilibrium, these balanced causal relationships will not be caught.

This basic point about taking state change data from dynamic processes has particular implications for CF-violations. For processes that re-equilibrate after 50 ms, for instance, a measurement device that samples the process at higher time scales, such as 500ms, will miss the re-equilibration. Thus, even though the system does not violate causal stability, it will behave as if it does, as it will appear that there is a conditional independence between two variables across some range of values, namely, the range between the initial state and the state to which the system was perturbed. In particular, if

we do not know what the time scale is, or is likely to be, for re-equilibration, we cannot ensure that a persisting probabilistic independence between two variables in question is genuine or a consequence of an overly fast re-equilibration timescale.

Not only does comparative time scales matter for apparent CF-violations; there are also possibilities for phase-matched cycles that that will make a non-CF-violating oscillating system appear to violate CF. Some systems develop equilibrium mechanisms that result in slight oscillations above and below a target state. If the measurements from this system are taken with a frequency that closely matches that of the rate of oscillation, then the measurements will pick out the same positions in the cycle, essentially rendering the oscillation invisible. This would constitute an apparent CF-violation as well.

Predicting possible CF-violations, real or apparent, requires information about the dynamic and evolved complexity of the systems in question, the particular equilibrium states they display, the time scale for re-establishment of equilibrium compared with the time scale of measurement, and/or the cycle length for cyclical processes.

6. Conclusion

To summarize briefly: some kinds of systems, especially those studied in the so-called ‘special sciences’, are likely to display the kinds of structural features that lead to CF-violations, such as mechanisms for equilibrium maintenance across a range of variable values. Some systems that do not have CF-violating DAGs may nevertheless generate CF-violating data sets. When we are considering the inferences made from probabilistic relationships in data to a DAG for the underlying system, and do not already have the DAG in hand, we cannot distinguish between genuine and merely apparent CF-

violations; both will cause the same epistemic difficulties for scientists, which is why merely apparent CF-violations deserve broader attention.

It's important to note that I am not discounting the extraordinary achievements in formal epistemology and causal modeling that have marked the last two decades of research on this topic. The steps forward in this field have been monumental, including the development of methods by which to reduce some of the issues arising from CF-violations (such as Zhang and Spirtes 2008). Rather, my goal is to clarify the ways in which apparent CF-violations can arise, the kinds of structural features a system might display that would increase the likelihood of CF-violation, and to bring this issue from discussion in formal epistemology into consideration of scientific practice more broadly.

References

- Berger, V.J., and A.D. Kharazova. 1997. "Mechanisms of Salinity Adaptations in Marine Mollusks." *Hydrobiologia* 355 (1-3): 115-126.
- Cartwright, Nancy. 1999. "Causal Diversity and the Markov Condition." *Synthese* 121 (1-2): 3-27.
- Cartwright, Nancy. 2002. "Against Modularity, the Causal Markov Condition, and Any Link between the Two: Comments on Hausman and Woodward." *The British Journal for the Philosophy of Science* 53 (3): 411-453.
- Cartwright, Nancy. 2006. "From Metaphysics to Method: Comments on Manipulability and the Causal Markov Condition." *The British Journal for the Philosophy of Science* 57(1): 197-218.
- Hausman, Daniel M. and James Woodward. 1999. "Independence, Invariance and the Causal Markov Condition." *The British Journal for the Philosophy of Science* 50 (4): 521-583.
- Hausman, Daniel M. and James Woodward. 2004. "Modularity and the Causal Markov Condition: A Restatement." *The British Journal for the Philosophy of Science* 55 (1): 147-161.
- Mitchell, Sandra D. 2003. *Biological Complexity and Integrative Pluralism*. Cambridge Studies in Philosophy and Biology: Cambridge University Press.
- Mitchell, Sandra D. 2008. "Exporting Causal Knowledge in Evolutionary and Developmental Biology." *Philosophy of Science* 75 (5): 697-706.
- Pearl, Judea. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press.

- Russo, Federica and Jon Williamson. 2007. "Interpreting Causality in the Health Sciences." *International Studies in the Philosophy of Science* 21 (2): 157-170.
- Steel, Daniel. 2006. "Indeterminism and the Causal Markov Condition." *The British Journal for the Philosophy of Science* 56 (1): 3-26.
- Spirtes, Peter, Clark Glymour, and Richard Scheines. 2000. *Causation, Prediction, and Search*. Cambridge, MA: The MIT Press.
- Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.
- Woodward, James. 2010. "Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation." *Biology and Philosophy*. 25 (3): 287-318.
- Zhang, Jiji and Peter Spirtes. 2008. "Detection of Unfaithfulness and Robust Causal Inference." *Minds and Machines* 18 (2): 239-271.