King's Research Portal

DOI:

[10.5840/jphil201611311](https://doi.org/10.5840/jphil201611311)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Wolff, J. E. (2016). Using Defaults to Understand Token Causation. *JOURNAL OF PHILOSOPHY*, 113(1), 5-26.
<https://doi.org/10.5840/jphil201611311>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Using defaults to understand token causation¹

J.E. Wolff

Published in *The Journal of Philosophy*, 113(1):5-26, 2016, DOI:
 10.5840/jphil201611311
 Please cite the published version

Abstract

Recent literature on causation invokes a distinction between deviant and default behavior to account for token causation. Critical examination of two prominent attempts to employ a distinction between deviants and defaults reveals that the distinction is far from clear. I clarify and develop the distinction by appeal to the notion of a modally robust process, and show how the distinction can be employed by causal process theorists to respond to cases of causation by omission. This shows that the default/deviant distinction is not so much a tool for counterfactual accounts of causation, but rather for causal process theory.

Introduction

With his strong left foot, the striker sends the ball towards the goal, the goalie already beaten on his right. Just before the goal line, one of the central defenders sticks out his foot, and the ball is knocked wide behind the goal for a corner kick. Disappointed we sit back down. We had seen that one in goal already. How do we

¹ I would like to thank Alexi Burgess, Douglas Cannon, Matthew Parrott, Justin Tiehen, Ariela Tubert, as well as two anonymous referees for extremely helpful comments on earlier versions of this paper. I would also like to acknowledge support from the University of Hong Kong and the Alexander von Humboldt Foundation.

know that a goal would have been scored, had it not been for the central defender's last minute action?

In this paper I pursue the idea that we anticipate the behavior of systems based on their default behavior. The default behavior of a well-struck soccer ball is to continue on its trajectory, not to swerve randomly, come to a sudden halt, or to explode. The ball's default behavior allows us to form expectations, and it also influences our causal reasoning. It was the defender's action that prevented the goal; it was his doing that caused the corner kick. A system's default behavior, intuitively, is the behavior a system exhibits under normal circumstances, when nothing interferes with it.

The question I will address is how we should understand this intuitive appeal to 'default' behavior, and how we can use a more careful distinction between default and deviant behavior in our accounts of token causation. Token causation is commonly taken to be a relation between two particular events c and e : the defender's sticking out his foot in that crucial moment in the 87th minute caused the ball's ending up behind the goal.

In section one I look at two ways in which a distinction between defaults and deviants has recently been employed to fix the counterfactual analysis of causation. Neither approach is entirely satisfactory, but they show how a distinction between defaults and deviants might be useful for understanding token causation. In section two I clarify the notion of a default by showing that defaults must be understood in terms of processes. Since this notion of a default initially turns out to be suitable primarily for fundamental physics, it is not immediately apparent how to employ it

for the cases of interest to theories of token causation. I then offer an analogy to show how we can expand the narrow notion of a default to typical cases of token causation, by showing what it means for a complex system to be in a default state. Using the clarified and expanded notion of a default from section two, I show in section three how a process theory of causation can employ this improved notion of a default to solve one of its most pressing problem cases, causation by omission, while remaining an adequate model for prevention and pre-emption.

1. Defaults and the counterfactual theory of causation

The distinction between defaults and deviants has recently come to the forefront in discussions of the counterfactual theory of causation. The counterfactual theory of causation begins with the simple idea that an event c is a cause of an event e iff e would not have occurred, had c not occurred, that is iff e counterfactually depends on c . As we all know, this simple analysis will not do. There are systematic ways to create counterexamples: either by showing that e does not counterfactually depend on c , despite the fact that we are inclined to count c as a (token) cause of e (as in cases of preemption and symmetric overdetermination), or by showing that e counterfactually depends on some event or condition g that we are not inclined to count as a cause of e (typical cases are background conditions and omissions). Counterexamples generated along these lines, as well as ever more sophisticated responses to them, make up the bulk of the literature on the counterfactual analysis

of causation.² Typical counterexamples to the idea that counterfactual dependence is sufficient for causation are cases of omission.

Omission

Assassin poisons Victim's coffee. Victim's Bodyguard is in possession of an antidote, but does not administer it. Victim drinks the coffee and dies shortly afterwards.

Victim's death after drinking the coffee counterfactually depends both on the poisoning and on the failure to administer the antidote. But while we certainly count the poisoning of the coffee as a cause of Victim's death, most of us tend to be less inclined to call the failure to administer the antidote a cause of death. Even if we are willing to count the failure to administer the antidote among the causes of victim's death, we are still inclined to say that there is a difference in the role played by the poisoning of the coffee and the failure to administer the antidote. The challenge for any counterfactual dependence account of causation is to explain the difference in the role of the two events.

Omission and similar cases where events do not seem to be caused by all of the events they counterfactually depend on, prompt Christopher Hitchcock (2007) to offer a "disjunctive" theory of causation; he thinks that for some cases, counterfactual dependence is necessary and sufficient for token causation, while in other cases it is either not necessary or not sufficient. Any such disjunctive theory obviously owes us an account of what distinguishes these different cases, otherwise the theory has no bite. Hitchcock provides a detailed recipe for telling whether the relation between two events is that of token causation.

² For an excellent overview, see Collins et. al. 2004.

First, Hitchcock says, build a causal model, where “[a] causal model is an ordered pair $\langle \mathbf{V}, \mathbf{E} \rangle$, where $\mathbf{V}$ is a set of variables and $\mathbf{E}$ is a set of equations among these variables.” (Hitchcock, 2007, 499) Then, crucially, ask whether this model is *self-contained*. A self-contained model, according to Hitchcock is a model in which “every variable $X \in \mathbf{V}$ satisfies PSR [the Principle of Sufficient Reason, see below]” (Hitchcock, 2007, 509). More specifically, ask not just whether the model is self-contained, but whether certain elements of the model, *causal networks*, are self-contained. A causal model can contain certain kinds of subsets of variables, causal networks, which, for the purpose of analyzing the causal relationships, may have to be distinguished.³

Whether a causal network is self-contained is the crucial question for the disjunctive theory of token causation. If the network is self-contained, counterfactual dependence is a necessary and sufficient condition for causation. Counterexamples to the counterfactual theory of causation arise only in cases where the causal networks are not self-contained. Hitchcock suggests that when there is counterfactual dependence between events in networks that are not self-contained, the dependence is ‘parasitic’, and not sufficient for token causation. On the other hand, one event may be a token cause of another event in a network that is not self-contained, even if there is no counterfactual dependence between them; to come to the right conclusions on those cases, more complicated theories are needed.

³ More formally: “Let $\langle \mathbf{V}, \mathbf{E} \rangle$ be a causal model, and let $X, Y \in \mathbf{V}$. The *causal network* connecting X to Y in $\langle \mathbf{V}, \mathbf{E} \rangle$ is the set $N \subset \mathbf{V}$ that contains exactly X, Y , and all variables $Z \in \mathbf{V}$ lying on a directed path from X to Y in $\langle \mathbf{V}, \mathbf{E} \rangle$ ” (Hitchcock, 2007, 509).

It is in distinguishing self-contained networks from other networks, that Hitchcock appeals to a distinction between defaults and deviants. Intuitively, in a self-contained network, the deviant value of any variable other than the first one can be explained by the deviant value of one or more of its parent variables. The idea behind this is the *Principle of Sufficient Reason*, which Hitchcock formulates in terms of default and deviant values of variables: “when a set of variables all take their default value, they cannot by themselves cause another variable to take a deviant value” (Hitchcock, 2007, 508).

These tools can be used to address the type of counterexample described above. With ‘0’ standing for the default value of a variable, and ‘1’ for the deviant value, the causal model for *Omission* looks like this:

Variables:

A=1 if Assassin poisons Victim’s coffee, 0 if not.

B=1 if Bodyguard administers the antidote, 0 if not.

D=1 if Victim dies, 0 if not.

Equations:

A=1

B=0

D= A & ¬B

Within this causal model, two causal networks can be distinguished: the causal network {A,D} and the causal network {B,D}. {A,D} is self-contained, and *D* depends counterfactually on *A*. Hence, given the analysis offered by Hitchcock, *A* counts as a token cause of *D*. But while *D* also counterfactually depends on *B*, the causal

network $\{B,D\}$ is not self-contained: D takes a deviant value if B takes a default value. By Hitchcock's analysis, that means the counterfactual dependence does not indicate a causal relationship. The difference between the two causal networks explains the difference in intuition between the two cases: while A is definitely a cause of D , in the case of B our intuitions are less clear.

Hitchcock's solution,⁴ then, depends on the distinction between self-contained and not self-contained networks, which is in turn defined in terms of the default/deviant distinction. Hitchcock's solution nicely explains our intuitions, in particular our inclination to distinguish the way in which Victim's death depends on the poisoning from the way in which it depends on Bodyguard's failure to administer the antidote. His recipe, however, is designed to handle only cases where counterfactual dependence is not sufficient. As he himself acknowledges, there is token causation outside of self-contained networks, and for those cases, more complicated theories are needed.

If Hitchcock's main concern has been with counterexamples to the idea that counterfactual dependence is sufficient for causation, Ned Hall (2007) is primarily interested in examples that seem to show that counterfactual dependence is not necessary for causation.⁵

⁴ Hitchcock offers similar explanations for other examples where counterfactual dependence does not seem to suffice for causation.

⁵ Another reason for introducing a distinction between default and deviant states are counterexamples to the structural equations account of causation. Hall constructs two scenarios about which we have differing causal intuitions—in one case we are inclined to

Preemption

Assassin puts a certain type of poison into victim's coffee. Unbeknownst to him, Badgirl also poisons Victim's coffee, but with a different poison that acts as an antidote to Assassin's poison, while still remaining lethal itself. Victim drinks the coffee and dies.

This type of example suggests that counterfactual dependence is not necessary for causation, since it seems that Badgirl's poisoning of the coffee is a cause of Victim's death, but it is not true that Victim would not have died had Badgirl not poisoned the coffee. After all, in that case Victim would have died from Assassin's poisoning instead. Hall suggests that in cases where we are inclined to judge that an event c is a cause of another event, e , despite the fact that e does not counterfactually depend on c , the counterfactual dependence is masked by surrounding events. According to Hall, what matters for causation is not just whether e counterfactually depends on c in the actual situation, but whether there is a nomologically possible situation in which it does.

The intuitive idea is clear: if it were not for Assassin's poisoning of the coffee, Victim's death would counterfactually depend on Badgirl's poisoning. So perhaps

count a particular event, c , as a cause of event e , in the other case we are not inclined to do so. Then he shows that the two scenarios have isomorphic structures, hence the structural equations account has to treat them alike. This prompts Hall to look for a non-structural difference between the two approaches, which he finds by distinguishing default from deviant states.

the reason we are inclined to call Badgirl's poisoning a cause of Victim's death is that we see the counterfactual dependence as hidden by the circumstances. In other circumstances, Victim's death would counterfactually depend on Badgirl's poisoning of the coffee. Of course, once we are willing to look for 'masked' or 'hidden' counterfactual dependence, we need to protect against a certain kind of paranoia—we need to have some limit in place that tells us when to look for masked dependence, otherwise we will start seeing counterfactual dependence even in cases where there is none. To prevent paranoia, we need another recipe.

Like Hitchcock, Hall begins by building a causal model. To reduce a complex situation, in which counterfactual dependence might be masked, to a simpler one, we construct a nomologically possible situation where only events occur that also occur in the original situation, but in which some previously deviant variables are now assigned default values.⁶ "In one situation, lots of events occur—that is, various bits of the world exhibit deviations from their default states. In another situation, strictly fewer events occur—that is, some of the bits of the world that are in deviant states in the first situation are in their default states instead; and every other bit is in the same state as it was." (Hall, 2007, 129) The only values we get to change by hand are the values of exogenous variables, that is, variables that have no parents in the model.

⁶ A similar account is offered by Stephen Yablo, who suggests that we 'hold certain things fixed' to find out whether an event 'de facto depends' on some other event (Yablo, 2002).

Applied to the case of *Preemption* we unmask the counterfactual dependence of Victim's death on Badgirl's poisoning the coffee, by reducing the original situation to a simpler version of it, where Assassin's actions are set to a default value of non-poisoning, and Badgirl's actions remain the same. In the reduced situation, Victim's death does counterfactually depend on Badgirl's poisoning the coffee, and hence we are back to a case where counterfactual dependence is necessary for causation.

Could the same argument be used to make Assassin's poisoning of the coffee a cause of victim's death? No, says Hall,⁷ because the situation where Badgirl does not poison the coffee is not a reduction of the original situation. For consider the state of Victim's coffee after Assassin has put in his poison. The default state for coffee is not to be poisoned, but since Assassin has put his poison in, that is not the state of the coffee. The actual state of the coffee—in the original situation—is poisoned, but with Badgirl's poison, not Assassin's, since Badgirl's poison neutralizes Assassin's. So a situation in which Assassin poisons the coffee, but Badgirl does not, is not a reduction of the original situation since it leaves one state in neither its default state, nor its actual state, but in a different deviant state.

Like Hitchcock, then, Hall uses a distinction between defaults and deviants to identify situations in which counterfactual dependence is an indicator for causal relationships, except that Hall looks at cases where counterfactual dependence does not seem to be necessary, whereas Hitchcock is interested in cases where it is not

⁷ Hall actually makes this point using an abstract 'neuron diagram'; the details of the example are mine.

sufficient. As Hitchcock (2009) has argued, however, Hall's approach relies on the background circumstances being deviants, since only then can they be 'set to defaults' to reveal masked dependence. If the background conditions that mask the dependence are themselves already defaults, no reduction of the situation is possible. Hall's approach, then, looks somewhat less promising than Hitchcock's.

2. Clarifying and expanding the notion of a default

Hall's and Hitchcock's discussions suggest that defaults are useful for understanding token causation. As we have seen, both of them use the term 'default' relatively loosely, applying it to states, the values of variables, and outcomes. Similarly, assigning the status of a default to some of these variables or states was usually done intuitively. While the specific assignments seem plausible enough, assigning default status in this manner can suggest that what counts as a default is largely a matter of context, and perhaps our own expectations.⁸

This may be fine for an analysis of how our causal reasoning works, but it is less clear that it will suffice for a metaphysical account of causation. If it turns out that our best causal reasoning relies on assigning the status of defaults to states or outcomes, then a metaphysical approach should say what the world must be like in

⁸ Hitchcock suggests that what counts as a default may depend on the "level of analysis" (2007, 506).

order for those assignments to be vindicated.⁹ My suggestion in this section is that assigning the status of defaults (or deviants) to states and outcomes requires processes. In the first instance, a default is a *default* relative to a process.¹⁰ In what follows I will use the term ‘behavior’ as a metaphysically neutral way of talking about the items which may be characterized as ‘defaults’, before developing a more technical analysis.¹¹ Intuitively, default behavior is uniform, normal, or orderly. The question is what we need to assume in order to make sense of these ascriptions.

Take one of the paradigms of default behavior—inertial motion.¹² Inertial motion goes well with the *Principle of Sufficient Reason*: if a system in inertial motion

⁹ Not everybody interested in understanding causal reasoning will be interested in pursuing this metaphysical project. Halpern (2008) seems primarily interested in causal reasoning without attempting to give a metaphysical account.

¹⁰ Hitchcock indeed at one point speaks of “deviant and default outcomes of various processes” (Hitchcock, 2007, 498), but he never suggests that states are defaults relative to processes.

¹¹ It is difficult to find a neutral term that fits all the relevant cases and examples. Although ‘behavior’ has its drawbacks, it seems to be a good fit for some of the central cases, like inertial motion. I hasten to add that while ‘behavior’ is often applied to the case of agents, I do not intend to imply that behavior must agential.

¹² Inertial motion is recognized by all parties to be default behavior; its status as such is especially emphasized by Tim Maudlin (2007) and Halpern and Hitchcock

changes its behavior, we have to look for a prior deviation of that system, or we have to look outside the system itself, to discover the reason for the change. Systems in inertial motion do not just change their behavior on their own. Notice, though, that in order to say, at any moment, that a body is in a state of inertial motion, we have to make reference to the earlier and later behavior of that body. It is only a state of inertial motion if the body is moving at uniform velocity, and whether the velocity at any given moment is uniform depends on what the velocity is immediately before and after.

To better describe what is going on in cases like these, I would like to borrow a conception of processes recently introduced by Helen Steward (Steward, 2013): the notion of a process as a modally robust individual. Proceeding from a distinction between events and processes, she suggests that we should recognize individual processes as entities with their own identity conditions, independent of the identity of their temporal parts. The motivating thought is that we sometimes recognize that a process, which was in fact interrupted, would have continued in a certain way, which suggests that we do not always treat processes as individuated by their (spatio-)temporal parts alone: the same processes could have gone on for longer, or it could have stopped earlier. To support such claims we need to think of a process as modally more robust than the mereological essentialist would allow; for

(2014), who write “The classic example is from Newtonian mechanics: the default behaviour of a body is to continue in a state of uniform motion” (p.21).

the mereological essentialist, any change in the temporal parts of a process constitutes a change in the identity of the process. Steward suggests that individual processes have “*modal robustness in virtue of form*” (Steward, 2013, 807).

This modal robustness in virtue of form is a property processes share with ordinary substances, like statues and cats. We usually think that a statue losing a finger, or a cat losing some hair on my couch, remain the self-same statue or cat, despite having undergone a change in their respective parts. Steward’s proposal is that we say the same about a process that has been interrupted, and is hence missing some of its temporal parts. In both cases it is possible to distinguish the process or substance from its spatial or temporal parts, because its identity depends on its form, and not merely on the parts that compose it. The form of a process, then, has to provide identity conditions for a process, which allow us to pick out the same process across different modal contexts. The question is, of course, which features of a process might be suitable to do so.

I am proposing that the notion of a default has its natural home in this conception of a process as an individual process. Such a process has temporal parts, and may consist of heterogeneous stages (Steward, 2013, 805), but these different stages hang together in a suitable way, given by the form of the process. I would like to call a stage of a process that sufficiently preserves the form of the process a default relative to that process. In an individual process there is a natural ‘next step’ or ‘next stage’ that would fit with the process as a whole, by preserving the form of the process. Such natural next stages, I suggest, are default stages, and since which

stage is a natural next step depends on the process, default stages will be defaults relative to processes.

To apply this general idea to the cases of token causation from which we started, we need to understand better the sources of 'form' for a process. Steward is primarily concerned with cases in the philosophy of action, and accordingly proposes agents' intentions as the source of form for processes. The cases I will be concerned with here are of a different sort, so I will be introducing different sources of form for processes. On the one hand, processes may be said to have a certain form in virtue of being uniform, as in the case of inertial motion. In those cases, the default next step will be to preserve the uniformity of the process by leaving certain quantities unchanged, in this case, the velocity. On the other hand, I will suggest that processes may be said to have a certain form in virtue of contributing to the stability of systems.

As long as our interest is only in Newtonian bodies and processes of inertial motion, the characterization of the processes, and accordingly, that of the default states, is going to be relatively simple. Indeed, because processes in fundamental physics seem to be most readily characterized as uniform or stable, causal process theorists have tried to build their account of causation on such processes, and have in turn characterized them in terms of conserved quantities Dowe (2000).¹³

¹³ This is not to say that all processes in fundamental physics exhibit uniformity, or that accounts in terms of conserved quantities are unproblematic in all areas of fundamental physics. But insofar as we understand a process as uniform, the

Most cases of token causation in the literature, however, do not deal with examples from fundamental physics, but with mundane cases involving entities like human beings dying from poison. Such entities undergo numerous different processes at any given moment: Victim, for example, might be at rest in an armchair, while drinking coffee, breathing, reading the news, and, involuntarily, ingesting the poison put in the coffee by Assassin—to list only a fraction of the processes Victim undergoes. In those cases, it seems, we start from relatively stable systems, like Victim, and events that happen to them, the causes of which will typically not be identified at the level of fundamental physical processes.¹⁴

So far, then, our best attempt at giving a clear notion of ‘default’ is to say a stage is a default in a process, where the process is characterized by a uniformity or stability over time characteristic of certain processes in fundamental physics. On the other hand, our intended application of the notion of a default is to situations involving complicated, but relatively stable, systems undergoing numerous processes at the same time. Is there any way of bridging the gap?

Fundamental physics gives us one understanding of an individual process: such a process is a uniform unfolding of stages. This account of the form of a process

relevant uniformity will most readily be found in certain processes described by fundamental physics.

¹⁴ Whether there is genuine causation at levels other than fundamental physics is of course a controversial question.

allows us to characterize defaults: a stage that retains the uniformity of the process is a default relative to that process. Call this the narrow notion of a default. But perhaps the uniformity retained does not have to be the uniformity of a process. Could it be the uniformity of a system as well?

A way of extending the notion of a default to cover cases beyond fundamental physics might then be this: states of complex systems can be more or less stable. Uniformity is a kind of stability, so perhaps default behavior preserves stability in general, not just the uniformity of fundamental processes. By analogy with processes, then, we can call the stable states of a complex system its default states. Call this the wider notion of a default. What the narrow and the wider notion of a default have in common is that default status is always held with respect to the stability of some other entity: a process or a system. It is relative to maintaining the stability of the other entity that a stage or a state can be said to be a default. In the case of complex systems, it is the systems that exhibit stability, not the processes.

Maintaining default states for a complex system typically requires the interaction of several processes, which themselves may or may not be uniform. So initially it would be difficult to understand the stages of those processes as defaults or deviants, since the processes lack uniformity and our best understanding of defaults relative to processes was as stages that retain the uniformity of the process. However, importing the extended notion of a default from above, we can say that a stage is a default relative to a process a complex system undergoes if it contributes to the stability of the system undergoing the process. The 'next' or default stage in a

process is the one that contributes to the stability of the system, even if it does not retain uniformity in the process.

This extension of the notion of a default to cover non-uniform processes fits nicely with the description of processes in biology as homeostatic. For example, our body's ability to retain a relatively stable internal temperature relies on a number of processes that are not themselves uniform in the way inertial motion is uniform, but which together result in a relatively stable state. In characterizing this state as a default state for us as organisms, we can then characterize stages in these processes as defaults or deviants, relative to whether or not they contribute to the stability of our body temperature. This allows us to apply the broader notion of a default to cover stages in processes beyond the processes of fundamental physics.¹⁵

I have now extended the notion of a default in two steps. First in extending the notion of a default from the stages of fundamental processes that are uniform to the stable states of complex systems, and then secondly in using this extended notion to define default stages for the numerous processes contributing to the stability of a complex system. It is unsurprising that we have to proceed by analogy here. After all, we are trying to extend a notion fit for one set of circumstances to fit a different set of circumstances. Crucially, at the end of the analogy we have found a

¹⁵ This also seems to fit well with Steward's idea that "many processes have natural developments—that there are certain types of unfoldings in the world which are structured in such a way that the termination point, product, or ongoing production cycle is the norm"(Steward, 2013, 807).

way of assigning the status of defaults to stages of processes even for cases involving the complicated systems of interest to token causation.

We can now say that defaults are either stages in processes that retain the uniformity of the process, or stages in processes that retain the stability of a complex system undergoing those processes. Since the processes are modally robust, we can say what stage would have been the default stage to occur in a given process, even if the process was in fact interrupted. Moreover, I have used the notion of a default state to characterize the stability of complex systems. A somewhat accidental feature of this approach is that causes for events like Victim's death, which disrupt the stability of the state of a complex system, will typically be found by looking at processes that involve only parts of that complex system. The reason is that complex systems are complex precisely because they require the working together of multiple processes, some of which may only occur in parts of that system. Can we use this refined understanding of defaults to solve problem cases in token causation?

3 Process theory, defaults, and omissions

Since defaults turn out to be best understood as being relative to processes, a natural place to turn is the process theory of causation. Defenders of the process theory of causation have long held that processes are fundamental to our understanding of causation. A causal process (as opposed to a pseudo-process) has recently been characterized as "a world line of an object that possesses a conserved quantity", and a causal interaction as "an intersection of world lines that involves exchange of a conserved quantity" (Dowe, 2000, 90). To show that defaults belong in

the toolbox of the process theorist, we not only need to show that defaults require processes, but also that defaults can do useful work for the process theorist.¹⁶ In this section I will show how the process theorist can employ defaults to address one of her most pressing problems: causation by omission.

Causation by omission is a problem for many theories of causation, but especially so for the process theorist. Causal processes are processes that transmit causal influence, and while different proponents of the view have suggested different criteria for identifying the features that allow processes to do so, all agree that transmission requires that the events are physically connected. It seems odd, however, to speak of transmission in cases where one event is a ‘negative’ event, that is, an ‘absence’ or an ‘omission’.¹⁷ If causal interactions occur at the intersection of two processes, and such an interaction involves physical connection, then how are we to understand cases where something happens because of an absence?

¹⁶ In what follows I take myself to be contributing to the project—called for by Hitchcock (2004)—of showing that causal processes and interactions can do real philosophical work; I am not trying to defend a particular version of the causal process account.

¹⁷ Regardless of their (alleged) role in causation, the question of what omissions are, if anything, is itself subject to metaphysical debate. For a detailed recent contribution see Clarke (2014), who argues that omissions are a special case of absence of action, but who aims to remain neutral with respect to the causal role of omissions.

The trouble is, as Jonathan Schaffer (2000) has argued, that we often want to count cases where absences or omissions are involved as cases of causation: the absence of vitamin C causes scurvy, the lack of oxygen in Victim's bloodstream caused Victim's death, and the power failure caused several people to be trapped in an elevator. All of these claims *prima facie* seem to count absences or negative events as token causes, which suggests that it is a severe disadvantage if a theory has to deny that absences can be causes.

There have been a variety of responses to the problem of causation by omission. Some philosophers, like Dowe (2001) and Hall (2004) have suggested that we need to distinguish two different concepts of causation, one of which is applicable only to cases of positive causation. Others, like Beebe (2004) and Varzi (2007) have rejected the idea that omissions play a genuine role in causation, and have instead suggested that alleged cases of causation by omission are instead to be understood (merely) as cases of causal explanations. The adequacy of a causal explanation may well depend on matters of moral responsibility, and it is with respect to such matters that omissions may become relevant.¹⁸ But this is quite different from the truth of such explanations, which only depends on positive

¹⁸ Clarke's view seems to go a step further, in that he defines omissions as absence of action in cases where "an action of the sort that is absent was called for by some norm, standard, or ideal" (Clarke, 2014, 33). For his narrow notion of omission, absences of action can only be understood as omissions against the backdrop of norms and standards.

events. While normative considerations seem to play a role in explaining our causal judgments for certain cases of omissions, Hitchcock and Knobe (2009) argue that normative considerations enter into causal explanations and judgments involving positive events as well. Hitchcock opts for a more radical response, suggesting that token causation in general is ‘not objective’, and that the locus of objectivity is causal structure, as represented by causal equations, not token causation (Halpern and Hitchcock, 2014). Accordingly Halpern and Hitchcock (2014) develop a notion of ‘defaults’ and ‘normalcy’, which tries to capture our actual, subjective judgments of token causation, in cases of omissions and elsewhere. I agree with Halpern and Hitchcock that causal structure is more fundamental than token causation, but I think their account falls short in two points.

First, to make the case clear we need a metaphysical account of causal structure—we know that structural equations are supposed to be the appropriate representation of the structure, but that does not yet tell us what that structure is. I suggest that the robust account of processes developed in the previous section can be employed here. Secondly, it seems a category mistake to relegate token causation to the merely subjective. Token causal judgments and explanations are indeed something we humans do, and are hence subject to our cognitive quirks, but token causation is that at which our judgments are aimed. It may very well be correct to think that token causation in fact has a secondary status compared to causal structure, but this secondary status should be expressed in terms of metaphysical relations, not by conflating token causation and token causal judgments. In what follows I intend to employ the notion of a default as developed in the previous

section to show both how token causation depends on processes, and secondly, how this allows us to understand omissions as genuine cases of causation, albeit ones of metaphysically secondary status.

Typical cases of token causation will be described as interactions of two processes, with the token effect lying at the intersection of the two processes. Crucially, the token effect will be a default with respect to one process, and a deviant with respect to the other process. When looking at a token effect, then, we need to ask two questions:

(1) What is the process relative to which e is a default stage? (2) What is the process relative to which e is a deviant stage? The cause of e will lie somewhere along the process relative to which e is a default. The intersections relevant to token causations are interruptions. A process p_1 interrupts another process p_2 if the outcome that is the intersection of p_1 and p_2 is a deviant relative to p_2 and a default relative to p_1 .

In *Poisoning*, Assassin poisons Victim's coffee. Victim drinks the coffee. Victim dies. We recognize Victim's death as a deviant outcome because staying alive is a relatively stable state for a healthy human being. It is not a stable state in the sense in which inertial motion is a uniform process, of course, since in order to maintain the stable state of being alive, many processes have to continue working together, whereas inertial motion is self-sufficient. Victim's death, then, requires the extended notion of a default, which we arrived at through the analogy laid out above.

The stable state of Victim's life was supported by several processes. To end that relatively stable state, one or more of these contributing processes must have

been interrupted. In searching for the cause of Victim's death, then, we look for a deviant relative to one of the processes supporting Victim's ongoing survival. A pathologist might determine that Victim's brain lacked oxygen for an extended period of time. The lack of oxygen in Victim's brain is a deviant stage relative to the process of respiration and transportation of oxygen in Victim's body. The cause of Victim's death will be found if we find the process that led to the interruption of the process transporting oxygen to Victim's brain. This process is of course the process of Victim's being poisoned: Victim's blood lacking oxygen due to important enzymes being blocked off, the presence of a suitable inhibitor, the ingestion of the inhibitor through Victim's drinking of the coffee, the presence of the inhibitor in the coffee due to Assassin's poisoning to coffee. Where we stop in tracing back this process is in part a matter of our interests. As pathologists, for example, we might be content to note that the cause of death was the presence of a suitable inhibitor in Victim's bloodstream, whereas as criminal investigators we will have to go beyond that to find out how the inhibitor got into Victim's bloodstream in the first place. The cause is not a single event, but a string of stages in an unfolding process.

My suggestion, then, is that we should understand token causation as the interruption of one process by another. What it means for one process to interrupt the other is in turn understood in terms of the default/deviant distinction: a process p_1 interrupts another process p_2 if the intersection of p_1 and p_2 is a default stage relative to p_1 and a deviant stage relative to p_2 .

Let's apply this recipe to *Omission* and *Prevention*. Hitchcock treats the two alike, but I actually think there are important differences between the two cases.

Omission, recall, was a case where Assassin poisons Victim's coffee, and Bodyguard, while in possession of the antidote, fails to administer it. As a result, Victim dies. *Prevention*, by contrast, is the case where Bodyguard in fact administers the antidote, and Victim survives. In *Omission*, Hitchcock argued, Assassin's poisoning of the coffee counts as a cause of Victim's death, because the poisoning and the death form a self-contained network, and the death is counterfactually dependent on the poisoning; the failure to administer the antidote, by contrast, does not count as a cause of Victim's death, although Victim's death is counterfactually dependent on it, because the failure to administer the antidote and the death do not form a self-contained network. Notice, though, that Hitchcock is forced to give the same analysis of *Prevention* (Hitchcock, 2007, 513). The causal model for *Prevention* is:

A=1 if Assassin poisons the coffee, 0 if not; B=1 if Bodyguard administers the antidote, 0 if not; D=1 if Victim dies, 0 if not. The equations are: $A=1$, $B=1$, $D=A \& \neg B$.

Once again the network $\{A,D\}$ is self-contained, which means that counterfactual dependence is necessary and sufficient for token causation. Since Victim's survival does not depend counterfactually on Assassin's poisoning of the coffee, Assassin's poisoning does not count as a cause of Victim's survival, which seems right. Also like in *Omission*, $\{B,D\}$ is not a self-contained network, and hence Bodyguard's action is also not a cause of Victim's survival, despite the fact that Victim's survival counterfactually depends on it, and for the same reason that Bodyguard's failure to administer an antidote in *Omission* was not a cause of Victim's death.¹⁹ Hitchcock's recipe rules that the dependence of Victim's survival

¹⁹ To be fair, Hitchcock strictly speaking only says that his recipe does not deliver a

on Bodyguard's administering the antidote in *Prevention* is 'parasitic' on Assassin's actions in the same way that the dependence of Victim's death on Bodyguard's inaction had been parasitic in *Omission*. But it seems to me at least, that there is a striking difference between the two cases: I am much more inclined to call Bodyguard's administering of the antidote a cause of Victim's survival, than I am to call his failure to do so in *Omission* a cause of Victim's death. Hitchcock's approach does not have the resources to respond to this intuition of difference.

Here is what my proposal says about *Omission*: Just as in *Poisoning*, Victim's death is a deviant outcome relative to the stable state of Victim's being alive. In looking for a cause of Victim's death, we find that one of the processes required for maintaining the stable state of Victim's life was interrupted by another process, which could be traced back all the way to Assassin poisoning the coffee. Bodyguard's action or inaction does not come into it, which is why we are not especially inclined to call Bodyguard's failure to administer the antidote a cause of Victim's death. But why does Bodyguard's inaction 'not come into it'?

Intuitively, because the relevant processes take place within Victim's body, whereas Bodyguard is a system (largely) independent of Victim. While I think that

clear verdict in cases where the network is not self-contained. Even so, his recipe still has to treat Bodyguard's inaction in *Omission* and Bodyguard's action in *Prevention* as causally alike, so either both are causes despite the fact that the networks they are in are not self-contained, or neither is. The intuition of difference remains unexplained.

captures something important, it is easy to find it dissatisfying. After all, Victim's body is by no means a closed system: without oxygen to breathe he would be dead without the poison. Moreover, presuming that Bodyguard is in the vicinity when all of this occurs, there are any number of physical processes involving both Victim and Bodyguard: light being reflected and absorbed, sounds being transmitted... . So what is different about the processes resulting in Victim's death?

Intuitively, again, the processes that do take place between Bodyguard and Victim are irrelevant to the poisoning. Using the analysis from above, we can say more precisely that the reason Bodyguard's actions are *irrelevant*, is that they are not interruptions of either the processes supporting Victim's survival, or of the process interrupting one of those processes. To be relevant to a particular case of token causation, a process, or a stage of a process, must either be (a part of) one of the processes the intersection of which constitutes the token effect, or it must be (part of) a process that is an interruption of such a process. None of the processes Bodyguard is engaged in are relevant in this sense. We may still be inclined to hold Bodyguard responsible in *some* sense, but we should not mistake that for taking his lack of action as a cause of Victim's death. Perhaps his inaction amounts to a breach of contract, perhaps it is negligence.

Contrast this with the case of *Prevention*. Relative to the various processes involved in Victim's bodily functions, Victim's being alive at time t_1 after drinking the coffee is a default state. Relative to having been poisoned at t , however, Victim's being alive at t_1 is a deviant state. We need to look for a process relative to which Victim's being alive at t_1 is a default. Victim's survival is a default relative to the process that,

working backwards again, consists of: the binding of the poison by another substance, the presence of a substance with a strong affinity for the poison, the injection of that substance into Victim's bloodstream by Bodyguard. Bodyguard's administering the antidote is a cause of Victim's survival, because it is part of a process relative to which Victim's being alive at t_1 is a default state. Unlike in the case of *Omission*, Bodyguard performs an action that is a stage in a process, which is one of the two processes that have Victim's survival at their intersection. The improved notion of a default, then, allows us to explain differences among cases that Hitchcock's recipe treated alike.

What about cases where we do want to count omissions or absences as causes? Arguably such a case can be constructed from Victim's death above. Victim's death is brought about, ultimately, by absence of oxygen in Victim's bloodstream. So the absence of something, oxygen, caused the death. That looks like a case of a negative event acting as a cause, which is thought to be in conflict with process theory. But based on what has been said so far, the process theorist has an answer. The reason lack of oxygen is a cause in this case is that the absence of oxygen in the bloodstream is a deviant stage relative to the process of breathing, which is one of the processes supporting Victim's being alive. Victim's mouth and lungs are expanding and contracting as usual, and oxygen is being taken in. But as the oxygen reaches the bloodstream, instead of being transported to Victim's brain, it is bound by the inhibitors present in Victim's bloodstream. The process of transporting oxygen to Victim's brain is interrupted, and Victim dies. It is because Victim's survival requires a process that very much involves positive events and presences

(in particular: the presence of oxygen) that interrupting this process leads to Victim's death. The interruption itself is a positive event: the inhibitor binding the oxygen is about as physical a connection as anyone can ask for.

Using the notion of a default, the process theorist can count *some* cases of absences as cases of causation, without having to count all of them as causes. That is as it should be, given that we only want to count some cases of omissions as causally relevant, but not others. In fact, I think the use of defaults is more than just a convenient tool for the process theorist in these cases. What the discussion above shows is that 'omissions' or 'absences' only enter into the causal nexus when they are interruptions of processes. An absence becomes causally relevant as such only when it is a deviant stage, and as I argued above, the status of being a default or a deviant is assigned relative to a process. It is important to note, though, that processes are interrupted by other processes. Absences depend metaphysically on presences, just like holes in an Emmental depend on the surrounding cheese.

The process theorist thereby gives an account of causation by omission, without committing to omissions as ontologically fundamental entities. It will be helpful to contrast this approach to omissions with a slightly less committed, more Humean account of (apparent) causation by omission. On this account, developed by Helen Beebe, omissions (and negative events more generally) do not exist and are never literally involved in causation. We give causal explanations involving omissions, but we have to distinguish the role omissions play in such causal

explanations from the role true causes play in such explanations. Explanations involving omissions can be adequate as explanations, because the omission is relevant, but omissions are not what make such explanations true. The relevance of omissions is a matter of norms, not causes.

An important difference between the account developed here and accounts like Beebee's is that I do not deny the existence of omissions or absences, I merely relegate them to metaphysically secondary status vis-à-vis processes. The difference lies in the idea that the process theorist, as I have characterized her, is able to say more about how omissions become relevant to *causation*, not just causal explanation. They are relevant when they constitute interruptions of processes. But their status as interruptions depends on the more fundamental existence of modally robust processes, of which we can say how they would have unfolded, had they not been interrupted. An interruption is a particular stage in a process which is deviant relevant to that process. At least some such deviant stages will be constituted by an absence or lack of something, instead of a presence.

Both Beebee's account and my account agree that causal explanations provide modal information, that is, information about how things would have gone under different circumstances. In the case of causal explanations invoking absences, for Beebee that means we are given information about the causal structure of the closest possible worlds in which the—actually absent—event indeed occurs. For example, to explain Victim's death by reference to the lack of oxygen in his blood is to describe how Victim's causal history would have gone, had oxygen been present in his blood: he would not have died. The modal information does not concern

absences, but presences in nearby possible worlds. By contrast, on the view I'm proposing here, the modal information is about what would have happened to a modally robust process which takes place in this world, and which is in fact interrupted. Process-theoretic explanations provide information about the causal structure of this world, and do not need to invoke the causal structure of other possible worlds. Of course, the price to pay for this is the acceptance of un-Humean, modally robust entities.

For some the price of accepting (modally robust) individual processes and stable systems may be too high. The alternative would be to go subjective on absences, not only in cases of omissions, but cases of negative events in general. There is nothing special about Victim's blood lacking oxygen, it's just that it is not what we expected. Relevance becomes a matter purely of our expectations, not a feature tied to anything going on in the world independent of such expectations. That is a possibility, of course, and it is closer to Hume's original account than contemporary Humean accounts. I do not find it particularly attractive. True, what counts as regular and stable is tied to human (time-)scales. But that does not mean that this stability is not real. It just means that like all stability, it is relative stability, not absolute stability.

What about *Preemption*? As we saw in the discussion of Hall's approach, preemptions are usually thought of as problems for the counterfactual theorist. What needs to be explained in Preemption cases is why an event, *c*, is a cause of an event *e*, but another event, *a*, is not. We arrive at this problem, because the

counterfactual analysis of causation seems to result in neither *a* nor *c* being causes of *e*, since if one had not happened, the other would still have brought about *e*. To use the example from above:

Preemption

Assassin puts a certain type of poison into Victim's coffee. Unbeknownst to him, Badgirl also poisons Victim's coffee, but with a different poison that acts as an antidote to Assassin's poison, while still remaining lethal itself. Victim drinks the coffee and dies.

For the process theorist, this case is not especially puzzling. Victim's death is a deviant outcome, since Victim's being alive was a relatively stable state for Victim to be in, just like in all the other examples before. Moreover, there is a process relative to which it is a default stage: the process that begins with Badgirl's poisoning of the coffee. So far, everything is completely analogous to the simple case of *Poisoning*. The situation in *Preemption* is slightly more complicated only in that there is a third process going on, the one that starts with Assassin's poisoning of the coffee. Badgirl's poisoning of the coffee first interrupts the process started by Assassin's poisoning, and later interrupts one of the processes supporting the stable state of Victim's life. The process started by Badgirl's action then interrupts two other processes, one of which had the same default outcome as the one started by Badgirl: Victim's death. But from the process theorist's perspective, the fact that another process was underway which had the same default outcome does nothing to undermine the status of the process started by Badgirl as the causally efficacious one. For the process theorist all that matters is which process goes

through, not what would have happened otherwise. Hence there is no need to mask the other process to reveal counterfactual dependence, as Hall suggested.

The causal process theorist, then, can use defaults to address both causation by omission and causation by preemption. That is a significant advantage over either Hitchcock's or Hall's recipe, since their respective recipes were designed to deal with only one type of counterexample. Moreover, the use of defaults suggested here was able to explain differences between cases that Hitchcock's account had to treat alike.

A final problem one might consider for the account I've described might be situations where the intersection of two processes results in the interruption of both.²⁰

Explosion

A block of dynamite is at rest on the ground. A flame on a fuse is traveling towards it. As the two processes intersect, an explosion occurs. As a result, the block of dynamite is no longer intact, and the flame has ceased to travel along the fuse. In this case, both processes are interrupted as a result of the intersection. My proposal as originally stated suggests that token causation happens where two processes intersect, and the intersection is a deviant with respect to one process, but a default with respect to the other. In the example as described, the intersection appears to be deviant with respect to both processes. I believe my original

²⁰ I would like to thank one anonymous referee for bringing such cases to my attention.

proposal can be generalized without much difficulty to cover cases like these, by weakening the requirement. Token causation is the intersection of two or more processes, which is a deviant stage with respect to at least one of the processes involved. This statement is weaker than the previous one, but also more generally, applying both to cases like *Explosion*, but also more generally to cases involving more than one process.

I have argued that the notion of a default is useful in understanding cases of token causation. I first showed how the notion of a default is currently employed by counterfactual theorists to respond to counterexamples. I then clarified the notion of a default, and it turned out that the status of being a default is relative to a process, either in the straightforward sense of contributing to the uniformity of a process, or in the extended sense of being a stage in a process that contributes to the stability of a complex system. In the final section I have used this improved notion of a default to show how the process theorist can respond to cases of causation by omission. I hope that what this discussion has shown is that the notion of a default is a useful tool for understanding token causation, but that it is most useful for the process theorist.

Even those unpersuaded by the alternative account of causation sketched in the final part of the paper may find some use in the notion of a default developed in this paper. For while Hitchcock and Hall introduced this notion in the context of the, generally Humean, idea that causal dependence should be understood as counterfactual dependence, what I have shown in this paper is that the notion of a default is best understood as belonging to a rather un-Humean conception of

causation involving modally robust individual processes. A committed Humean will likely find such a notion unattractive. If so, she should also stay away from the notion of defaults. Friends of the process theory, by contrast, may find the notion helpful in further developing their view.

References

- Helen Beebe, "Causing and Nothingness", in John Collins, Ned Hall, and L.A. Paul, eds., *Causation and Counterfactuals* (Cambridge, MA: MIT Press, 2004), pp. 291–308
- Randolph Clarke, *Omissions: Agency, Metaphysics, and Responsibility* (Oxford: Oxford University Press, 2014)
- John Collins, Ned Hall, and L.A. Paul, eds., *Causation and Counterfactuals* (Cambridge, MA: MIT Press, 2004)
- Phil Dowe, *Physical Causation* (Cambridge: Cambridge University Press, 2000)
- Phil Dowe, "A counterfactual theory of prevention and 'causation' by omission" *Australasian Journal of Philosophy*, 79(2) (2001):216–226.
- Ned Hall, "Two concepts of causation", in John Collins, Ned Hall, and L.A. Paul, eds., *Causation and Counterfactuals* (Cambridge, MA: MIT Press, 2004), pp. 225–277
- Ned Hall, "Structural equations and causation", *Philosophical Studies*, 132, (2007):109–136
- Joseph Y. Halpern, "Defaults and normality in causal structures", *Proceedings, Eleventh International Conference on Principles of Knowledge, Representation, and Reasoning*, 2008

Joseph Y. Halpern and Christopher Hitchcock, "Graded causation and defaults", *The British Journal for the Philosophy of Science*, Advance Access (doi:10.1093/bjps/axt050) (2014):1–45

Christopher Hitchcock, *Causal processes and interactions: What are they good for?* *Philosophy of Science*, 71(5), (2004):932–941

Christopher Hitchcock, "Prevention, Preemption, and the Principle of Sufficient Reason", *Philosophical Review*, 106(4), (2007): 495–531

Christopher Hitchcock, "Structural equations and causation: Six counterexamples", *Philosophical Studies*, 144 (2009): 391–401.

Christopher Hitchcock and Joshua Knobe, "Cause and Norm", *Journal of Philosophy*, 106(11), (2009): 587–612

Tim Maudlin, "Causation, Counterfactuals, and the Third Factor", in *The Metaphysics Within Physics*, (Oxford: Oxford University Press: 2007), Chapter 5

Jonathan Schaffer, "Causation by disconnection", *Philosophy of Science*, 67, (2000): 285– 300

Helen Steward, "Processes, continuants, and individuals", *Mind*, 122(487), (2013): 781– 812

Achille C. Varzi, "Omission and causal explanation", in F. Castellani and J. Quitterer, eds., *Agency and Causation in the Human Sciences* (Paderborn: Mentis Verlag, 2007) pp. 155–167

Stephen Yablo, "De facto dependence", *The Journal of Philosophy*, 99(3), (2002): 130–148