




2011

Singular and General Causal Relations: A Mechanist Perspective

Stuart Glennan

Butler University, sglennan@butler.edu

Follow this and additional works at: http://digitalcommons.butler.edu/facsch_papers

 Part of the [Epistemology Commons](#), [Philosophy of Mind Commons](#), [Philosophy of Science Commons](#), and the [Social and Behavioral Sciences Commons](#)

Recommended Citation

Singular and General Causal Relations: A Mechanist Perspective (2011). In *Causality in the Sciences*, edited by Phyllis McKay Illari, Renee Russo and Jon Williamson. Oxford: Oxford University Press, 789-817.

This Article is brought to you for free and open access by the College of Liberal Arts & Sciences at Digital Commons @ Butler University. It has been accepted for inclusion in Scholarship and Professional Work - LAS by an authorized administrator of Digital Commons @ Butler University. For more information, please contact fgaede@butler.edu.

*Singular and General Causal Relations: A Mechanist Perspective**

Stuart Glennan

Butler University

July 15, 2009

1. Introduction

Causal claims can be divided into two kinds – singular (token) and general (type). For instance, we claim generally that heating butter causes it to melt and, regarding a single case, that the heat that I applied to the butter in my kitchen yesterday evening caused it to melt. One of the central puzzles about the nature of causation concerns the relationship between these sorts of claims and the causal relations that underlie them. There seem to be two main options. What I shall call the generalist view holds that singular causal relations obtain because they are instances of a general causal regularity or law. What makes it true that the heat I applied to the butter last night caused it to melt is that there is a general causal law that whenever butter is heated past a certain point it will melt. What I shall call the singularist view holds the converse position. Singular causal relations can obtain even if they are not instances of causal regularities or laws, and what makes causal generalizations true, when they are true, is that they correctly describe a pattern of singular instances of causally related events. It is a basic fact that last night my heating the butter caused it to melt, and the general claim that heating butter causes it to melt is true only because it happens that in most or all of the individual cases, heating butter causes it to melt. To put the matter succinctly, the generalist holds that general causal relations make singular causal claims true, while the singularist holds that singular causal relations that make general causal claims true.¹

* Forthcoming in *Causality in the Sciences*, edited by Phyllis McKay Illari, Renee Russo and Jon Williamson. Oxford: Oxford University Press

¹ Two other papers which explicitly discuss the relationship between singular and general causal claims are (Hitchcock 1995) and (Russo and Williamson 2009). What I call singularist and generalist

My aim in this paper is to make a case for the singularist view from the perspective of a mechanical theory of causation (Glennan 1996, 1997, 2010, forthcoming), and to explain what, from this perspective, causal generalizations mean, and what role they play within the mechanical theory. Prior to making this argument, it is important to clarify the relationship between the singularist/generalist distinction and another distinction widely discussed in the contemporary literature on causality. It is now commonly held that there are two concepts of cause—or at least that our causal assertions make two different sorts of claims (Hall 2004, Hitchcock 2007, Godfrey-Smith 2010, Glennan forthcoming). On the one hand, causes are said to produce or bring about effects. On the other hand, causes are said to depend upon, be relevant to, or make a difference to their effects. The case for thinking of these concepts as distinct is that there appear to be instances in which something can be a cause in one of these senses and fail to be a cause in the other. Two phenomena that illustrate this point are overdetermination and causation by omission. A paradigmatic example of overdetermination is a prisoner being executed by firing squad. In such a case a particular soldier's shot produces a wound that causes death, but the soldier's shot does make a difference to the prisoner's death because the other soldiers' shots were each sufficient to cause the death. Had the first soldier's shot not hit the prisoner, the prisoner still would have died. In cases of causation by omission, on the other hand, we appear to have causes that make a difference but that are not productive. Suppose I rear-end a car because I fail to brake. In such a case my failure to brake clearly made a difference to the occurrence of the accident. Had I braked, I would not have hit the car. But my omission cannot be said to have produced the collision, because my omission is not, properly speaking, an event or occurrence. What produced the collision was the forward momentum of my car, which was produced by my earlier pressing of the accelerator.

In what follows, I shall refer to these two concepts of cause as productivity and relevance. The significance of the productivity/relevance distinction for our discussion of singular and general causes is that the singularist view seems to fit more naturally with the productivity

approaches, Russo and Williamson call bottom-up and top-down strategies and Hitchcock calls Humean (or neo-Humean) and generalization approaches. A third option holds that the truth of singular and general causal claims are independent of each other. This position has been argued for principally by Eells (1991). I shall not discuss this position directly, but I hope to argue indirectly against it by providing a singularist account that addresses the concerns that motivate it .

approach, while the generalist view seems to more naturally make sense of relevance or difference-making. The connection between these two distinctions is evident in an early paper by Sober(1984) that makes a case for there being two concepts of cause. Sober's two concepts of cause are not causal productivity and causal relevance but are rather token causation and property (or type) causation. Nonetheless, Sober makes the case that token and type causation are distinct by appealing to the difference between productivity and relevance. Sober does this in the context of probabilistic theories of causality. Most probabilistic theories of causality are type-level difference-making accounts. Probabilistic theories of causality typically assert that causes must raise the probabilities of their effects. The standard problem for difference-making accounts is that there appear to be singular causal processes in which events that lower the probability of an outcome nonetheless are productively connected to the effect – sometimes called the problem of “doing it the hard way.” Many of these examples involve golf balls and squirrels. Suppose I have hit a putt that is heading cleanly toward the hole and there is a high probability of the ball going in. As the ball roles towards the hole, a squirrel runs up and kicks it off its path, but fortuitously the ball ricochets off an acorn that has just dropped to the green and bounces into the hole. Sober argues that the correct way to analyze this case is to argue that “the kind of kick” the squirrel made is a type-level negative cause of holing putts, in the sense that such kicks are negatively probabilistically relevant to holing putts, but that the particular squirrel kick is a token cause of holing this particular putt, because the particular process by which the putt made it to the hole “traces back” to the squirrel's kick.

Notwithstanding Sober's argument, causal relevance is not essentially a property or type-level notion. Difference-making can be understood in the single case by appeal to counterfactuals. The case of omission discussed above is an example of just such a case. When I say that my failure to brake was causally relevant to my rear-ending the car, I am talking about an omission that made a difference in a particular case. But while intuitive appeal to counterfactuals allows us to make sense of singular causal relevance claims, it remains to be seen whether we can understand these singular counterfactuals in a way that does not implicitly make reference to general causal claims. The central claim of this paper is that the mechanical theory will provide a way to meet this challenge.

In the next section of the paper we will review some traditional arguments for and against singularism. The remainder of the paper is concerned with describing the relationship between three sorts of theories of causation: process theories, mechanical theories and manipulability theories. In section three we will see how process theories provide a singularist account of causal relations, but in so doing fail to capture relations of causal relevance. In section four, I introduce mechanical theories, seeing how they contrast with process theories, particularly in regard to their treatment of causal relevance. In section five, I examine the relationship between mechanical theories and manipulability theories. The upshot of this examination is that the two sorts of theories should not be thought of as competitors but as describing different aspects of the nature of causal relation. In the final section of the paper I make the case that the mechanical theory as it has been developed in the previous sections really supports a singularist metaphysics.

2. Preliminary Arguments for and Against Singularism

The modern singularist view begins as a reaction to Hume's regularity theory of causation. Hume's belief that all knowledge of matters of fact derives ultimately from impressions, combined with his view that we have no impression of a necessary connection in a single case, leads to his view that singular causal claims are true because they are instances of regular patterns of association. So, for instance, billiard ball *a*'s striking billiard ball *b* causes billiard ball *b* to move, because (1) billiard ball *a* did strike billiard ball *b*, (2) billiard ball *b* did start to move, and (3) *ceteris paribus*, whenever one billiard ball strikes another, the second begins to move. The motivation for his regularity theory is essentially epistemological and pragmatic. Singular causal sequences are instances of regular causal sequences, because, unless a singular causal sequence is an instance of a regular causal sequence, it would be impossible to recognize it as causal. Moreover, if cause-effect sequences are instances of regularities, then it is possible to predict and control effects by observing or manipulating their regular causes.

Two of the most widely discussed singularist critiques of Hume come from Ducasse ([1926] 1993) and Anscombe ([1971] 1993). Both reject the epistemological strictures that suggest it is not possible to observe causal relations in the single case. Ducasse argues that the problem arises from Hume's presumption that the connection between causally related events was some third entity, analogous to the *relata*. He writes:

Hume's view that no connection between a cause and its effect is objectively observable would be correct only under the assumption that a 'connection' is an entity of the same sort as the terms themselves between which it holds, that is, for Hume and his followers, a sense impression. ... [But] the fact is the causal connection is not a sensation at all, but a relation. ... We observe it whenever we perceive that a certain change is the only one to have taken a place immediately before, in the immediate environment of another (Ducasse [1926] 1993).

In describing causality as a relation, Ducasse is suggesting that it is a relation like contiguity. When we observe two people sitting beside each other, we do not observe the first person, the second person and some third thing – "besideness", but rather we just observe that the first person is beside the second person. Similarly, Ducasse argues, when one event causes a second, we observe the first event, the second event, and the fact that they are causally related.

Anscombe's argument is Wittgensteinian in character. People learn to use a variety of specific causal concepts – her examples are "*scrape, push, wet, carry, eat, burn, knock over, keep off, squash, make* (e.g., noises, paper boats), *hurt*" (Anscombe [1971] 1993: 93). Only when one has mastered specific causal concepts is one able to master the highly general concept of cause. Moreover, mastery of causal concepts is essentially connected with mastery of the use of words describing objects and events:

If we care to imagine languages in which no special causal concepts are represented, then no description of the use of a word in such languages will be able to present it as meaning cause. Nor will it even contain words for natural kinds of stuff, nor yet words equivalent to 'body', 'wind', or 'fire'. For learning to use special causal verbs is part and parcel of learning to apply the concepts answering to these and many other substantives. As surely as we learned to call people by name or report from seeing it that a cat was on the table, we also learned to report from having observed it that someone drank up the milk or that the dog made a funny noise or that things were cut or broken by whatever we saw cut or break them (ibid.).

This picture is Wittgensteinian in the sense that it eschews what Wittgenstein in the *Philosophical Investigations* called the Augustinian picture, in which words are tags that are hung on things that one "observes". For Anscombe, like for Wittgenstein, acquisition of a concept involves mastery of certain techniques. These techniques are parts of language games in which a variety of words – nouns, verbs and other kinds – are connected to behaviors and social practices. There is no "observation" apart from mastery of these techniques, and once such mastery is achieved, observation of causal connections is no more or less problematic than observations of the objects that enter into them.

These arguments about how causal knowledge is acquired are helpful to the singularist's case but not decisive. The singularist's position is a metaphysical rather than an epistemological one, and the possibility of acquiring singular causal knowledge is neither necessary nor sufficient for establishing this metaphysical position. Davidson's position, for instance, is that Ducasse was correct that it was possible to know that a singular causal relation obtained without knowledge of a general causal law, but that the fact that a singular causal relation obtained entailed that there exists some law, even if we do not know what it is ([1967] 1993: 84-5).

Both Anscombe and Ducasse believe however that there are metaphysical or conceptual grounds for singularism. Anscombe's singularism stems from an observation she finds

so obvious as to seem trite.... Causality consists in the derivedness of an effect from its causes. This is the core, the common feature, of causality in its various kinds. Effects derive from, arise out of, come of, their causes. For example, everyone will grant that physical parenthood is a causal relation. Here the derivation is material, by fission. Now analysis in terms of necessity or universality does not tell us of this derivedness of the effect; rather it forgets about that (Anscombe [1971] 1993: 92)

Anscombe here articulates the productive conception of cause. The fetus was produced by the interaction of one egg and one sperm, and the arrival of the baby in the world was produced by the act of labor. It may be the case that this is how all babies are conceived and most babies come into the world, but these general facts need not be true for the singular causal claims to hold. Greek myths tell us that the goddess Athena had a most unusual birth – springing fully armored from Zeus' head after Hephaestus cracked it open with an axe. The story is doubtless false, but there is nothing inconceivable in such a singular birth; the fact that in general whacking heads with axes is not a way to produce children does not entail that the story of Athena's birth is wrong. Ducasse echoes this point about the relation between singular causal claims and causal generalizations:

... [T]he cause of a particular event [is defined] in terms of but a single occurrence of it, and thus in no way involves the supposition that it, or one like it, ever has occurred before or ever will again. The supposition of recurrence is thus wholly irrelevant to the meaning of cause; that supposition is relevant only to the meaning of law. And recurrence becomes related at all to causation only when a law is considered which happens to be a generalization of facts themselves individually causal to begin with (Ducasse [1926] 1993: 129)

While Ducasse and Anscombe have made a strong case for the singularist intuition, in my view neither of them has provided an adequate explanation of just what this singular causal relationship is. Ducasse does offer a reductive definition of cause, but, for reasons I shall not

explore, it seems wholly inadequate. Anscombe does not attempt to define what she means by determination. Presumably this is because she feels that this relation is both unanalysable and directly observable. While I do not doubt that in an ordinary sense we are quite capable of observing causal relationships, I do not think Anscombe's conclusion will do. While our conception of causality may originate in our typically successful observations of ordinary things pushing and scraping, identifying causal relationships can be far more complex. In the first place, there are circumstances – like magic shows – where our observations of ordinary causal relationships can be quite off the mark. Secondly, we often make causal inferences without observing a causal relation – as when we infer from a patient's symptoms that they have interacted with an infectious agent. An advocate of a singularist approach must then say something more about the nature of the causal relation.

While the arguments of Anscombe and Ducasse provide *prima facie* grounds for doubting Hume's view and for adopting a singularist perspective, the singularist perspective faces some important difficulties. Hitchcock (1995), discusses one difficulty, which concerns the semantic relationship between singular and general causal claim. What Hitchcock calls the generalization strategy supposes that singular causal claims are basic, and that general causal claims should be analyzed as generalizations over these singular causal claims. For instance, the general causal claim that smoking causes cancer is true because it generalizes over true singular causal claims – that Emily's smoking caused her to get cancer, that Edward's smoking cause him to get cancer, and so on. The problem with this strategy is there is not a simple relationship between generalizations and their instances. In probabilistic causal generalizations it is too strong to suppose that every instance of the singular causal relation obtain. Notwithstanding the general causal connection, not all who smoke get lung cancer. Perhaps one might treat a generalization like this as an existential one, but Hitchcock argues that there are true causal generalizations that have *no* instances. His example (ibid.: 265) is that eating one kilogram of uranium 235 causes death. This generalization, he claims, "is true in virtues of certain features of human physiology and the physics of nuclear chain reactions; however, no one has ever died in this unusual way and it is unlikely that anyone ever will." (ibid.). Hitchcock's example cannot be analyzed either as:

(x)(x's ingesting 1kg of U235 causes x to die)

or

$(\exists x)(x\text{'s ingesting 1kg of U235 causes } x \text{ to die})$

Both of these claims are false because no one has ingested 1kg of U235.

I would argue here that while Hitchcock's objection undercuts a natural sort of analysis of the relationship between singular and general causal claims, it doesn't undermine singularism as such. The problem here though doesn't have to do with singularism but with the fact that general causal claims have counterfactual import. Hitchcock's example is plausibly analyzed as:

$(x)(\text{if } x \text{ were to ingest 1kg of U235, } x\text{'s ingesting 1kg of U235 would cause } x \text{ to die})$

The singularist can maintain that the truth of this general counterfactual claim would depend upon the truth of singular counterfactual claims – for instance, that if Emily were to ingest 1kg of U235 her ingesting it would cause her to die.

Russo and Williamson (2009) raise an epistemological objection to singularism, arguing that the singularist (or as they call it, the bottom-up) causal metaphysics is hard to square with actual practices of inference in the sciences. They divide the sorts of evidence available into evidence of (singular) mechanisms and evidence of (general) difference-making, and claim that in the health sciences at least to establish a cause one must have evidence of both types. Their argument is based upon the analysis of causal inference in the case of autopsy:

To determine that *Alf's heart attack was a cause of his death*, the medical practitioner needs to have evidence both that there is a viable biological mechanism linking heart-attack and death *and* that the heart attack made a difference to his death. ... At the generic level, in order to establish that *pneumonia is a cause of death in hospital patients*, those conducting an academic autopsy need to be aware of evidence *both* of a mechanism linking pneumonia and death, *and* that pneumonia makes a significant difference to death in the population in question. ... The proponent of the mechanistic analysis cannot explain why, in cases where there is excellent mechanistic evidence, evidence of difference-making is also required.

Contrary to Russo and Williamson, I think the mechanistic analysis can easily explain why evidence of difference-making is required. In their first case, the "viable biological mechanism linking heart attack and death" is a generic description of a mechanism. Not all heart attacks cause death. Whether one does depends upon the details of the heart attack, the state of the victim's other vital systems and the place and circumstances of the heart attack. The fact that heart attacks on some occasions are linked via a physiological mechanism to a person's death makes a heart attack a *prima facie* candidate for the cause of a particular death. But to establish

the heart attack as a cause of death in Alf's case, one would have to show that it made a difference in this case. One would have to show that had Alf not had the heart attack, he would not have died. If, for instance, Alf was suffering from sepsis and the sepsis brought about a failure of a number of organs including the heart, and that, given all these conditions, the heart attack did not make a difference to the death.² Similar arguments could be offered in response to other cases to show that a diversity of epistemic methods does not undermine the metaphysical position of the singularist.

3. Process Theories and Singularism

Probably the most prominent attempt to provide a positive account of the nature of singular causal relations involves what I call *process theories*.³ Process theories assert that a cause is related to an effect via a series of processes and interactions. Processes are world-lines of objects that propagate causal influence through space-time, while interactions involve intersections of these world-lines in which properties of the processes are changed. Here is a simple example: Gretchen's throwing a baseball causes the window to break because the motion of her throw (an interaction of the ball and arm processes) leads to the flight of the ball (a process) that leads to the impact with the window (an interaction between the ball process and the window process) that produces the break. Note here that the process in question is a particular process involving particular objects at a particular place in space and time.

Process theories emerge historically as a response to difficulties with probabilistic theories of causality. Salmon (1980) suggests such a theory is required as a response to probability-lowering cause cases like the squirrel/golf ball scenario, and Sober's (1984) suggestions about token causation are also in keeping with this view. The view has been most thoroughly

² Identifications of cause of death would seem to involve a host of pragmatic factors, including social and legal conventions. How does one choose between proximal and distal causes, between environmental or internal causes, between various overdetermining causes? I am inclined to think that there is no objective answer to the question of what is "the cause of death."

³ Process theories are sometimes called transference accounts (cf. Craver 2007) or causal-mechanical accounts, or physical accounts. See (Dowe 2000, 2008, 2010) for a more nuanced discussion of theories in this tradition.

developed in the work of Wesley Salmon (1984, 1994) and, more recently, in Dowe's conserved quantity theory (2000). Process theorists agree that causal influence is transmitted via processes and interactions but differ on how properly to characterize processes and interactions. For the sake of definiteness, I will focus on Dowe's theory, though the objections I raise apply generally to process theories.

On Dowe's conserved quantity theory, causal processes are world lines of objects that possess conserved quantities, while causal interactions are intersections of causal processes in which conserved quantities are exchanged (Dowe 2000: 90). Conserved quantities are things like mass-energy, linear momentum and charge (ibid.: 91). Dowe characterizes his theory as a theory of physical causation, because it seeks to define causal production, propagation and interaction by appeal to our best physical theories. His is what he calls an empirical rather than a conceptual analysis of cause. While a conceptual analysis tells us about the meaning and use of causal concepts in ordinary language, an empirical analysis attempts to tell us what causation is in the world, apart from our ways of talking about it. The idea is that our causal talk picks out various instances of causal relations between events, and that we can look to physical theory to find out what physical characteristics mark the relations between these causes and effects.

While process theories offer the promise of yielding a theory of singular causation, they appear to fall victim to a series of objections involving causal relevance. The first class of objections work by identifying sequences of processes and interactions that connect a putative cause to an effect but which are nonetheless causally irrelevant to an effect. This sort of objection was originally raised by Hitchcock as a criticism of Salmon's theory, but it has been widely discussed with a number of variations (Hitchcock 1995, Salmon 1997, Glennan 2002, Craver 2007). A pool player chalks a cue stick, strikes the cue ball with the stick, the cue ball strikes the eight ball, which drops into the pocket. The motions of the stick and balls are paradigm causal processes, and the interaction of these processes involves exchanges of conserved quantities – namely momentum. The problem raised by Hitchcock is that it might be the case that a cue stick transfers a chalk mark from the stick to the ball and from the first ball to the second. These transfers are interactions, presumably involving exchanges of conserved quantities. The problem is that the transference of the chalk mark is causally irrelevant to the outcome of the shot, so the chalking of the cue is not a cause of the sinking of the eight ball. Examples like this can be multiplied indefinitely and point to the fact that the existence of a

continuous sequence of causes and interactions linking two events is not a sufficient condition for the first event being a cause of the second.

While this first objection suggests that the existence of an a continuous sequence of intermediate processes and interactions is not a sufficient condition for causal connection, omission and prevention problems suggest that the condition is not necessary. In cases of omission, it is the non-occurrence of some potential preventing cause that causes (or at least allows) some effect to occur. For instance, my failure to turn off the alarm when I walked in the door caused the police to come to my house. In cases of prevention, the occurrence of some event prevents another event from occurring. For instance, my catching the vase as it topples off the shelf prevents it from breaking on the floor. The problem with omission and prevention is that either the putative cause (in omission) or the putative effect (in prevention) is a non-occurrence. These non-occurrences are problematic for process theories because there can be no set of processes that link non-occurring omissions to effects or preventive events to non-occurring effects. This second class of objections, like the first, suggests that the process theory fails to capture causal relevance, non-occurrences cannot be causes or effects in process accounts, even though non-occurrences can be causally relevant to effect events in omission or cause events can be causally relevant to non-occurring effects in prevention.

A third important objection concerns what might be called the reductionist character of process theories. Process theories are typically theories of physical causation, in the sense that they seek to identify properties of causal connection in terms of concepts drawn from current physics. But the great majority of causal claims made both in ordinary and scientific discourse involve events and processes not described in the language of physics. We seem to have good evidence for the truth of causal claims in biology, psychology, economics, history, etc. that do not make any reference to the exchange of conserved quantities or any such concept drawn from physical theory. In fact, there is a long history arguing for the autonomy of these higher-level causal claims from physical theory (e.g. Fodor 1974, Kitcher 1984). These arguments

suggest that the “gory details” at the physical level are irrelevant to the truth of causal-explanatory claims in the higher-level sciences.⁴

4. The Mechanical Approach and the Process Approach

Process theories were developed largely as a response to problems that probabilistic and other type-level theories of causation had accounting for the productive character of singular causal relations. They addressed these problems, but at the cost of introducing problems with causal relevance. At this point one might reasonably embrace causal pluralism and say that we need different kinds of theories for different kinds of causes. I am sympathetic to this pluralism to a point. There are a variety of different sort of causal concepts that have a variety of uses within common-sense and scientific discourse, but this does not mean that there are multiple kinds of causality in the world. A metaphysical account should ideally provide a reasonably unified account of the nature of causation that can plausibly be connected with this variety of causal concepts. In this section, I argue that a mechanistic approach to causation can provide this unified account, and in particular that it can make sense both of relations of productivity and relevance.

While the mechanical approach to causation is quite different from the process approach, readers would be excused for thinking that processes and mechanisms come down to very much the same thing. In my first paper on the subject (Glennan 1996), I argued that, roughly speaking, two events were causally connected just in case there was an intervening mechanism. This sounds very much like the process theory. Matters are not helped by the fact that some process theorists have characterized their approach as mechanistic. Salmon (1984), for instance, calls his approach to explanation “causal-mechanical.”

⁴ Dowe (2000, 2008, 2010) has discussed these objections to process theories and has offered some correctives. Probably most importantly, Dowe insists that a theory of physical causation is just one aspect of a multi-faceted analysis of causation. So, for instance, he argues that omissions aren’t truly causes in the sense of physical causes, but he identifies an alternate relation of quasi-causation that attempts to capture intuitions about causal relevance. Dowe’s pluralism allows him to argue that his conserved quantity theory provides a correct empirical analysis of what is going on physically in causal processes, without requiring that theory to answer all questions about conceptual, empirical and explanatory aspects of causal relations.

The difference between the process theory and the mechanical theory lies in their rather different conceptions of what a mechanism is. For the process theorist, a mechanism just is a process of the sort described by their theory. To the mechanical theorist, however, a mechanism is a *system*. To get a sense of what this distinction amounts to, consider two widely discussed attempts to characterize a mechanism.

Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions (Machamer, Darden, and Craver 2000: 3) .

and

A mechanism for a behavior is a complex system that produces that behavior by the interaction of a number of parts, where the interactions between parts can be characterized by direct, invariant, change-relating generalizations (Glennan 1996: S344).

Machamer, Darden and Craver argue that mechanisms are *organized* and that they are productive of *regular changes*. Glennan argues that mechanisms are *systems* of interacting parts, where these interactions are characterized by *generalizations*. In both of these cases we see that mechanisms are systems that have a certain degree of stability. If we consider, for instance, the circulatory mechanism in vertebrates, this system contains a number of parts – heart, arteries, capillaries, veins, and blood – that are stable in their organization and operate in a regular way over the lifetime of an organism. A second important and shared feature of these conceptions of mechanisms is their hierarchical character. The parts that comprise a mechanism may themselves be complex systems whose capacities and dispositions are explained by the regular operations of the parts' own parts. Within the circulatory system, a part of that system – say the heart – will have parts (valves, chambers, and so forth) and these parts will have parts, and the parts at each level will have characteristic activities and interactions that are productive of the behavior of the mechanism of which they are a part. On this systems conception of mechanism, causal processes are understood as instances of the operation of mechanical systems. The circulation of blood through a particular animal's body is a continual process that results from the operation of the circulatory system.⁵

⁵ It is not plausible to suppose that all causal processes involve regular operations of a mechanical system of this sort. The squirrel who kicks the ball against the acorn into the hole is a clear example of such an irregular causal process. Processes of this sort are examples of what I call ephemeral mechanisms (Glennan under review). A complete mechanistic account of causation needs to explain what these

Both the regular and hierarchical nature of this approach to mechanisms contrasts with the view of process theories. The process theory focuses on a single process at a single point in space and time, as in our example of Gretchen throwing the ball and breaking the window. The collection of entities – Gretchen, the ball and the window – do not in any ordinary sense form a system. They do not act in a regular way to produce a repeatable behavior. The Gretchen-breaking-the-window process is also not hierarchical, because the properties that are required to establish that Gretchen’s throw caused the window to break are basic physical properties – exchanges of conserved quantities like momentum. It is the regular and hierarchical nature of mechanisms that provide the resources to address relevance problems. Consider first the example of the chalked cue stick. If we treat the situation in terms of the account of mechanisms given in (Glennan 2002) we consider the chalk, the cue stick, the player, the two balls, the table and the pocket as parts of a system describe the direct invariant change-relating generalizations that describe interactions between these parts. Some of the generalizations involved in this case would like be in the form of equations describing the change in momentum a part as a function of the momentum of a previous part at the moment of their impact. Other generalizations would describe the effect of the table on the ball as it rolled along the table and was slowed by friction. The chalking would not be part of the description of the system because changes in the chalking would not produce changes in the motion of the ball.⁶

The hierarchical character of the mechanical approach is important in avoiding the objection leveled against process accounts that the account of causation is overly reductive. Unlike the process theories, which seek to identify a physical criterion like exchange of conserved quantities that characterizes all physical interactions, on the hierarchical mechanical approach, causal interactions can occur at multiple levels of organization. In a circulatory system for instance, one characterizes an interaction between the blood (as a fluid) and the heart as an

ephemeral connections are, how they are related to mechanisms on the systems conception, and how one can provide an account of productivity and relevance for connections mediated by such mechanisms.

⁶ We should note here that there is a simplification involved. I am no pool player, but I suppose that if chalking is anything more than a ritual affectation, it does have an effect on the trajectories of balls – presumably by making the cue surface less slippery and allowing the player to impart spin to cue balls. So, on a careful analysis, chalking may be relevant. What couldn’t be relevant is any inadvertent coloring of the ball.

interaction between parts that can be characterized in terms of change-relating generalizations describing the relationship between variables such as heart rate, blood pressure and rate of blood flow. What makes these relationships causal is that they can be described by these invariant generalizations of physiology. While there is a further mechanistic explanation of why these generalizations are true, the behavior of the mechanism will be largely invariant with respect to changes in the structure of blood and tissue at the cellular and sub-cellular level.

Craver has argued that my version of the mechanistic approach does not in fact provide a suitable solution to the problem of causal relevance. He makes his case by providing a description of a particular mechanism that characterizes that mechanism in terms of a set of causally irrelevant properties. The mechanism in question is the mechanism of long-term potentiation (LTP)— a mechanism for strengthening the connection between pre- and post-synaptic neurons by rapidly stimulating pre-synaptic neurons (Craver 2007: 92) . Craver offers a “bizarre description” of this mechanism:

A glutamate molecule with molecular weight w crosses the synaptic cleft at velocity v , collides with a passing protein, alters the position of amino acids in the NMDA receptor, and lowers the concentration of Na^+ in the intracellular fluid.

He goes on:

This description includes a set of parts and mechanistically explicable interactions. Each stage is linked via a mechanism to its predecessor. Yet no one would claim this is a good explanation of LTP. This is because the putative explanation is composed of irrelevant features of the synapse. It is not the molecular weight of the glutamate molecule or its velocity that matter, but rather its conformation and charge configuration... (ibid.: 92).

Craver’s strategy is just an application of Hitchcock’s chalked ball argument to a neuroscientific example. He claims that my approach cannot rule out this bizarre description, but I think this is incorrect. What it means to say that a feature like the velocity of a particular glutamate molecule is irrelevant is to say (counterfactually) that if the velocity of the molecule had been different, the mechanism would still have produced the same behavior, and to say (actually) that the mechanism, which involves a large number of these molecules that will move across the cleft at a variety of velocities, will produce the behavior it does in spite of these variations – both among molecules in the synapse at a particular time and between molecules traveling across the synapse at different times. The mechanism that produces long-term potentiation will utilize interactions that can be characterized by invariant change-relating generalizations. Change-relating generalizations describe functional relations between two or more variables

where an ideal intervention on one variable will bring about a change in another variable. According to Craver's "correct" description of the LTP mechanism for instance, the binding of glutamate to an NMDA receptor changes the conformation of the receptor in order to open a channel for Ca^{2+} (ibid.: 70). This is an interaction between glutamate molecules and the receptor in which one change – binding to the receptor – produces another change – opening a channel. The reason that irrelevant characteristics, like the velocity of the glutamate model, are not including in a description of the interactions is that the behavior of the mechanisms (and the generalizations describing interactions between its parts) are invariant under interventions that change these characteristics.

What appears to have happened here is that in his criticism Craver has appealed implicitly to the Salmon-Dowe conception of an interaction, as opposed to the sort that I advocate, in which the interactions must be interactions are part of a mechanism that produce a particular behavior and must interact in accordance with invariant change-relating generalizations. While the version of the mechanistic account I favor addresses the causal relevance problem, potential objections to the mechanical theory remain. First of all, it may be, because of the centrality of my appeal to invariant change-relating generalizations in characterizing what constitutes a mechanism, that the mechanistic theory is in reality just a version of Woodward's counterfactual-based manipulability theory. Second, it may also appear that the prominent appeal to generalizations solves the relevance problem only by moving away from the singularist stance. I shall address the first of these concerns in section five of this paper and the second in section six⁷

5. The Mechanical Approach and the Manipulability Approach

⁷For more on how the mechanistic theory addresses the relevance problem, see (Glennan forthcoming). That paper also discusses how the mechanistic approach handles cases of omission and prevention.

There has been considerable discussion in the literature about the relationship between the mechanistic approach and counterfactual approaches to causation. In (Glennan 1996) I argued that the mechanistic approach explained the truth conditions for counterfactuals in a way that was more epistemically and scientifically helpful than that of Lewis (1973); but in (Glennan 2002) I suggested that Woodward's counterfactual account of invariant generalizations was essential to characterizing an interaction between parts of a mechanism. Machamer and Bogen have argued that a mechanistic approach to causation allows one to avoid counterfactuals altogether, a point that has been criticized by Psillos, Woodward and myself (Bogen 2004, Glennan 2010, Psillos 2004, Woodward 2004). Woodward has suggested that one can give a definition of a mechanism, or at least a mechanical model, in terms of manipulability criteria(Woodward 2002). For myself, I have come to believe that the mechanical theory and the manipulability theory – at least as it is advocated by Woodward (2003), Pearl(2000), and Spirtes, Glymour and Scheines (2000)– are not really rival theories, but rather highlight different features of a unified conception of causation. The manipulability account emphasizes procedures for discovery, prediction and control. The mechanical account provides different sorts of resources for discovery and prediction, a metaphysical underpinning for the manipulability approach and an enriched conception of causal explanation.

To understand how the manipulability theory and the mechanical theory are related, it is necessary to examine how causal relationships are typically represented in the manipulability theory. I base my exposition on (Pearl 2000) and (Woodward 2003), which I take to represent different pieces of a single approach to causation⁸ Pearl and Woodward assume that causal relationships can be represented as relations between variables, where variables can either take Boolean values representing the occurrence or non-occurrence of an event or continuous quantities representing different values of variable magnitudes (like temperature, salinity, or velocity). Causal relationships between variables are represented by causal models (or structural equation models). Causal models consist of a set of exogenous variables, a set of endogenous variables, and for each endogenous variable, a function from some subset of the

⁸ Woodward (2003: 38) acknowledges his debt to Pearl. As Woodward sees it, his theory is complementary to Pearl's. While Pearl is more focused on questions of inference, Woodward is more concerned with providing explicit definitions of notions like being a total or contributing cause, and more generally with relating Pearl's approach to causation into the philosophical mainstream.

variables (endogenous or exogenous) – its parents – to that variable. (Pearl 2000: 203) A causal model will determine a directed acyclic graph (DAG). For instance, Figure 1 shows a DAG involving two exogenous and three endogenous variables:

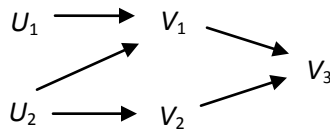


Figure 1: a directed acyclic graph

This graph would be determined by a causal model involving functions of the following variables:

$$v_1 = f_1(u_1, u_2)$$

$$v_2 = f_2(u_2)$$

$$v_3 = f_3(v_1, v_2)$$

Causal models provide a way of factoring dependence relationships among variables, so that the value of a variable depends only upon some subset of other variables – its Markovian parents. Parents are clearly represented in the DAG notation. We can see, for instance, that V_3 's parents are V_1 and V_2 , and that conditional on its parents, the probability of V_3 will be independent of U_1 and U_2 .

I claim that a causal model is a representation of a mechanism in the sense described in the mechanical theory. I will argue for this by way of an example. In (Glennan 1996), one of the examples I used of a mechanism was a toilet. A toilet is a mechanism for a certain behavior, which for purposes of illustration we can describe as follows: When the handle is pulled, water is released from the storage tank into the bowl, and the storage tank is refilled. Here is how the mechanism works. Pulling the handle (H) pulls a chain (C) which opens the flapper valve at the bottom of the tank (B_1). The open flapper valve allows the water to empty out of the tank (T_1), which has two effects. First, it causes the bottom valve to close (B_2) and second it causes a float to drop (F_1). The dropping of the float opens the float valve (V_1) allowing water to enter the tank. The opening of the float valve together with the closing of the flapper valve causes the tank to fill (T_2) which causes the float to rise (F_2) and the float valve to close (V_2). Here's the DAG:

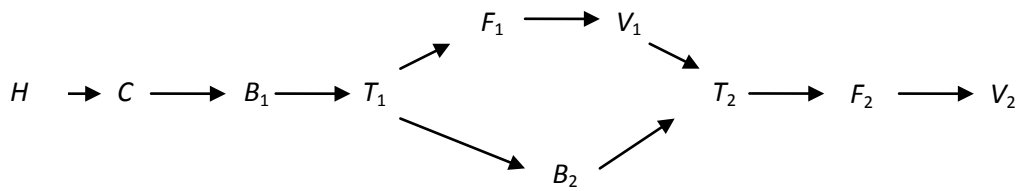


Figure 2: A DAG for a toilet mechanism

set of direct invariant change-relating generalizations. These generalizations will be in the form of equations characterizing the relation between each endogenous variable and its parent or parents (the f s). This representation does lose some information one might want in the model of a mechanism. First, the model characterizes the mechanism with a set of binary-valued variables. In fact the system contains some continuously varying magnitudes – the level of the water, the amount that the valve is opened, and so on. The causal modeling approach can however handle quantitative variables. Perhaps more importantly, the variables here do not represent parts (like the float valve) but rather changes in the state of the parts (like the opening or closing of the float valve). This is connected to the fact that the DAG representation does not illustrate the cyclical nature of the toilet mechanism, in which a part like a valve begins in a closed position, is opened, and is closed again.

If a causal model is to be a representation of a mechanism, it requires another feature, which Woodward calls *modularity*. If a model is modular, it must be possible in principle to intervene in order to change the value of a dependent variable without altering any of the other functional relationships in the model. Modularity so defined is a property of models, but it corresponds to an important property of the mechanisms modeled. In the case of the toilet, for instance, the modularity condition implies that one should be able to intervene on the state of one part and not thereby alter any of the functional relationships downstream of that part. For instance, one ought to be able to intervene on the chain, pulling it up, and not thereby interfere with the functional relationship between the water level and the float. Woodward offers the following motivation for the modularity requirement:

It is natural to suppose that if a system of equations correctly and fully represents the causal structure of some system, then those equations should be modular. One way of motivating this claim appeals to the idea that each equation in the system should represent the operation of a distinct causal mechanism. (Correlatively, each complete

set of arrows directed into each variable in a directed graph should also correspond to a distinct mechanism.) If we make the additional plausible assumption that a necessary condition for two mechanisms to be distinct is that it be possible in principle to interfere with the operation of one without interfering with the operation of the other and vice versa, we have a justification for requiring that systems of equations that fully and correctly represent causal structure should be modular. (Woodward 2003: 48)

While I have suggested that the causal model as a whole is, given the modularity assumption, a model of a mechanism, here Woodward (cf. Pearl 2000: secs. 1.3, 7.2.4) suggests that a single structural equation, representing the causal relationship between a node and its parent nodes represents a “distinct causal mechanism.” So, continuing with our example, an equation describing how pulling the chain relates to opening the flapper valve represents a distinct causal mechanism. Thus, it is clear in this instance that what Woodward and Pearl mean by a causal mechanism is what I have called an interaction between parts.

In another article, where Woodward responds directly to the systems account of mechanism he understands mechanisms more in the manner of (Glennan 1996) and (Machamer, Darden and Craver 2000). He offers the following description of a necessary condition for a causal model to be a model of a mechanism:

(MECH) a necessary condition for a representation to be an acceptable model of a mechanism is that the representation (i) describe an organized or structured set of parts or components, where (ii) the behavior of each component is described by a generalization that is invariant under interventions, and where (iii) the generalizations governing each component are also independently changeable, and where (iv) the representation allows us to see how, in virtue of (i), (ii) and (iii), the overall output of the mechanism will vary under manipulation of the input to each component and changes in the components themselves (Woodward 2002: S375).

In this characterization, the whole causal model (or directed graph) is the representation of a mechanism. My supposition is that when Woodward says that the behavior of each part is characterized by an invariant generalization, he really means that there are generalizations (perhaps multiple) describing both activities (behavior of the part) and interactions (relations between the behavior of one part and the behavior of directly connected parts). Given this definition, his view of what constitutes a mechanical model essentially coincides with (Glennan 1992). His view is similar to that of Machamer, Darden and Craver, except that he insists, contrary to Machamer (2004) and Bogen (2004) that the characterization of activities and interactions between parts of mechanisms requires counterfactuals.

It is significant that Woodward says both that causal models as a whole and equations representing the relation between parts of a mechanical system are both representations of mechanisms. This suggests that Woodward is in agreement with the account I sketched earlier regarding the hierarchical character of mechanisms, and specifically with a thesis I have defended elsewhere (Glennan 1996, 1997) that the generalizations describing the interactions between parts of mechanisms are in most cases mechanically explicable. Consider again the toilet mechanism. The chain connecting the lever to the flapper valve is a part of the toilet mechanism, but it is also itself a mechanism. The chain has parts (links) whose properties and interactions explain the behavior of the chain as a whole. Thus, an invariant change-relating generalization describing how pulling on the top of the chain will change the position of the bottom of the chain, will be a mechanically explicable generalization. Similarly, we can treat each link of the chain as a part of the chain and we can describe how each link is connected via a mechanically explicable change-relating generalization.

A final important similarity between the mechanical and manipulationist approach concerns the understanding of the semantics and epistemology of counterfactuals. In (Glennan 1996) I argue that one of the virtues of the mechanical approach over Lewis' counterfactual approach is that it provides an unproblematic way to understand and evaluate counterfactuals by reference to mechanisms. Given a model of a mechanism that exhibits the functional dependence of variables that represent the mechanism's parts and their properties, one evaluates a counterfactual claim by using the model to calculate what would happen if one were to intervene and fix the value of a variable to the antecedent of the counterfactual. For instance, in the case of the toilet one knows that if the chain were broken then the tank would not empty, because, if one were to intervene and break the link between *C* (the chain being pulled) and *B1* (the flapper valve opening), then all the events downstream of *C* would not occur.

Judea Pearl (2000: ch. 7) has developed a complete analysis of what he calls "structure-based counterfactuals" that formalizes this approach in terms of structural equation models and his "do operator." Like Glennan (1996), Pearl sees this analysis of providing an analysis of the truth conditions of counterfactuals that does not rely on the metaphysically extravagant and cognitively/epistemologically problematic possible-worlds semantics of Lewis:

In contrast with Lewis' theory, counterfactuals are not based on an abstract notion of similarity among hypothetical worlds; instead they rest directly on the mechanisms (or "laws" to be fancy) that produce those worlds and on the invariant properties of those mechanisms. Lewis' elusive "miracles" are replaced by principled minisurgeries, $do(X = x)$ which represent the minimal change (to a model) necessary for establishing the antecedent $X = x$ (for all [values of exogenous variables] u) (Pearl 2000: 239)⁹

It is revealing that Pearl so casually equates mechanisms and laws. It suggests that Pearl accepts the view argued for in (Glennan 1996, 1997) that laws are mechanically explicable.

These observations about the relationship between my analysis of causality and the manipulationist counterfactual approach of Woodward and Pearl should suffice to show how the two approaches are interconnected. The mechanical approach relies on the counterfactual approach because there is no way to define interactions between parts of mechanisms except by appeal to counterfactual-supporting generalizations. The counterfactual approach relies on the mechanical approach because the truth-conditions for counterfactuals depend upon the structure of mechanisms.

Stathis Psillos (2004) has also argued for an account of causation that seeks to "harmonize" mechanisms and counterfactuals, and most of what I have said here is consistent with Psillos' explication of the relation between mechanisms and counterfactuals. Psillos has, however, argued that "there is a sense in which the counterfactual approach is more basic than the mechanistic one in that a proper account of mechanisms depends on counterfactuals while counterfactuals need not be supported (or depend on) mechanisms" (Psillos 2004, 288). To complete our analysis of the relationship between the mechanical and manipulationist approaches, we need to assess Psillos' claim.

⁹ The formal analysis is not without difficulties. In particular, it may strike readers as problematic that the truth of a counterfactual claim is relativized to a model of a mechanism rather than to the mechanism itself. Of course if we are to appeal to mechanisms to make judgments about the truth of counterfactuals, we must inevitably rely on our models of these mechanisms, but we would like the truth itself not to depend on our representation. A second issue has to do with background conditions. For Pearl, a counterfactual "If it were the case that $X=x$, then it would be the case that $Y=y$ " will be true only if the model calculates $Y=y$ for all values of background variables U . This may be too strict. For instance, consider the counterfactual "If the flapper valve weren't to close properly, water would keep running into the bowl." Intuitively, this seems to be true, but its truth depends upon a certain background condition remaining constant, namely that the water supply to the toilet is kept on. If the water supply were included as an exogenous variable in the model, then the counterfactual would not be true. These complications do not seem to me to undermine the structural approach, but rather to be an inevitable consequence of the vagueness of counterfactuals.

The feature of my account upon which Psillos bases his claim is my claim that *most but not all* causal generalizations (or laws) are mechanically explicable. According to my account, the relationship between causes, causal generalizations and mechanisms is this: Two events are causally connected when there is an intervening mechanism. An intervening mechanism consists of a number of interacting parts, and these interactions are truly interactions (as opposed to accidental correlations) because they are described by invariant change-relating generalizations, which support counterfactuals. But if these generalizations are mechanically explicable, then what ultimately makes it true that the parts interact is that these interactions are produced by the operation of further, lower-level mechanisms. These mechanisms will in turn be systems of parts interacting in accordance with invariant change-relating generalizations, and these generalizations may too be mechanically explicable. Ultimately, however, one will reach a level where the parts of a mechanism interact, but where there is no further mechanism that explains this interaction. These are the fundamental interactions. What makes it the case that these relationships are truly interactions? The answer would seem to be that there is some basic, mechanically inexplicable, counterfactual dependence between events, perhaps one that holds in virtue of a fundamental law. As Psillos sees it, “the presence of a mechanism is *part* of a metaphysically sufficient condition for the truth of certain counterfactuals; the fully sufficient condition includes some facts about the fundamental laws that, ultimately, govern the behavior of the mechanism” (ibid.: 310).

Psillos is correct that the mechanical approach cannot eliminate counterfactuals, and because of this that it cannot provide a complete and reductive analysis of causal relations. This fact does not, however, entail Psillos’ asymmetry claim. In the first place, it appears that counterfactuals really do need mechanisms. We have seen in the above analysis of the manipulability theory that causal models are models of mechanisms and that in Pearl’s analysis of the semantics of counterfactuals, their truth conditions depend upon the structure of mechanisms. At least on Pearl’s analysis, and arguably on Woodward’s, what makes a certain counterfactual claim true is that there is a mechanism that would respond in a certain way to a manipulation. And like the mechanical theory, the manipulability approach faces a charge of *prima facie* circularity and uses the same strategy to respond to that circularity. Woodward, as he himself notes (2003: 103-107), defines causal relations in terms of the outcomes of possible interventions, and interventions are themselves kinds of causing. Woodward’s response to this potential objection is to argue that the circularity is not vicious. To determine whether two

variables X and Y are causally related, we must make know something about other causal relationships (e.g., between an intervention I and X), but not about the causal relationship between X and Y we are seeking to establish. True enough, but how do we know that an intervention I causes a change in X ? Presumably we know this because there is a mechanism connecting I and X . This will involve further variables (representing further parts and interactions), and how can we know that these variables are connected? By further interventions of course! Just as the mechanist must ultimately run out of nested mechanisms, Woodward must ultimately get to an intervention that cannot be further analyzed in terms of interventions on further mechanisms. One is left with the brute claim that if one were to intervene on X a change in Y would result. But this is not an analysis of the counterfactual dependence of X and Y ; it is just a restatement of it. Thus I would contend that the truth of causal claims according to the manipulability theory will depend upon an unanalyzed notion of counterfactual dependence.

Psillos' argument for the asymmetry of mechanisms and counterfactuals is really an epistemological one (Psillos 2004: 315-317). It is possible to construct a perfectly randomized experiment that establishes a causal connection between a treatment and a control. One can establish this connection without having any idea of what the mechanism is, and indeed, one could establish this even if there is no mechanism, but just a brute pattern of causal dependence. But one should not let this fact mislead one into thinking that the manipulationist approach has provided a metaphysical grounding for causal relations any more than the mechanistic account has. I do not think Woodward would object to this characterization of the situation, because he is emphatic that his analysis of causation is not reductive. It does not seek to ground the truth of causal claims in some ultimately non-causal state of affairs, but rather to explain the relationship between certain causal claims and others. This seems like a wise idea, especially given Woodward's focus on causal explanation, but, like the mechanical theory, Woodward's theory leaves crucial metaphysical questions unanswered. These have to do with the ultimate truth grounds for claims of counterfactual dependence at the level of fundamental physics, where the notion of causal interaction cannot be explicated by appeal to further mechanisms. How we understand these truth grounds will turn out to have a crucial impact on our understanding of our original question – whether causal claims are ultimately singular.

6. *The Mechanical Approach and the grounding of singular causal claims*

We are now finally in a position to make the argument that the mechanical approach supports a singularist view of causation. The basic reason why the mechanical approach is a singularist one is that it suggests that causal interactions are mediated by mechanisms, and mechanisms are particular systems of interacting parts, where these interactions occur at a particular place and time. On this view, causal generalizations are generalizations about the behavior of mechanisms, and they are true because mechanisms do or would behave in the way described on actual or hypothetical occasions. The problem that remains is that our definition of mechanism frequently makes reference to causal generalizations, and the suspicion will arise that the truth of singular causal claims depends ultimately on the truth of these generalizations, especially the non-mechanically explicable generalizations upon which the causal productivity of mechanisms would ultimately seem to depend.

Because the central issue concerns the implications of the hierarchical character of mechanisms for the status of singular causal relations, it will be helpful to have an abstract representation of a hierarchical mechanism, as in figure 3:

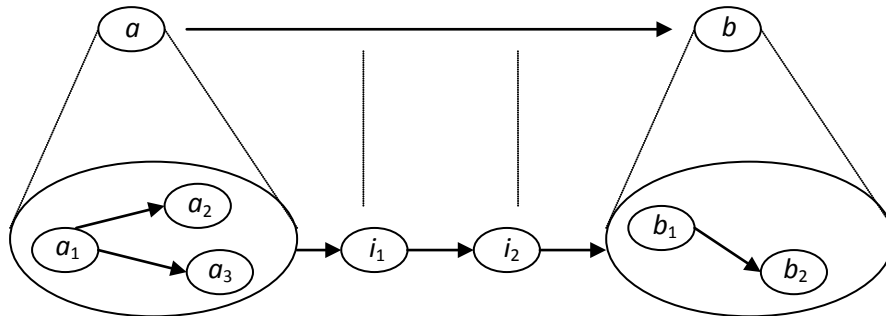


Figure 3: A hierarchical mechanism

The top of the diagram is a representation of two parts of a mechanism interacting. The dotted lines represent constitutive relations and the solid arrows represent causal interactions. The constitutive lines indicate that the part *a* is itself a mechanism with three parts and part *b* is a mechanism with two parts, where these parts interact as indicated by the arrows. The interaction between *a* and *b* may involve the operation of a further mechanism, as indicated by

the parts i_1 and i_2 . Although not pictured in this diagram, we should imagine that the parts of in the lower level of the diagram themselves have parts, and that the arrows representing interactions between these parts are themselves complex mechanisms with subparts. This hierarchy of mechanisms can go a long way down but will ultimately bottom out with fundamental parts and interactions. Imagine the parts to be atoms or corpuscles, much in the way Descartes imagined, and that these parts have some set of basic properties (e.g., mass, charge), interacting with each other in a manner determined by these properties.¹⁰

To defend the singularist interpretation, we must explicate the role of causal generalizations in this picture. The definition of a mechanism in (Glennan 2002) claims that interactions between parts “can be characterized by direct, invariant, change-relating generalizations.” What exactly are these generalizations and what is their relation to the interaction? One reading would be to treat these generalizations as a form of causal law, and to argue that the causal interactions are *governed* by these causal laws. If this were the case, the singular causal claim, that a change in a property of a produces a change in a property of b , would be made true by the causal law. This reading would undermine the singularist account, but it is not consistent with the hierarchical character of mechanisms and the mechanical explicability of these generalizations. The hierarchical picture suggests a second reading in which change-relating generalizations are statements that “characterize” the interaction, but the interaction itself involves the operation of the underlying mechanism and is not governed by the generalization. For example, there might be a change-relating generalization indicating that when I ingest caffeine, my motor activity increases. This generalization is true and reliable, but it simply characterizes the outcome of a complex metabolic mechanism.

¹⁰ A problem that I can only allude to here has to do with the implications of quantum mechanics for how we understand fundamental interactions. The picture I offer of a fundamental interaction is essentially a classical Cartesian/Newtonian one. Indeterminism in quantum mechanics raises some problems for this picture. More significant though are problems raised by the measurement problem and by violations of locality. The mechanistic picture seems to require bottom-level interactions that are local and have definite properties independent of measurement. I don’t have anything constructive to say about this problem. I can only offer as consolation the fact that except under special conditions quantum mechanical peculiarities wash out as one gets past sub-atomic scales. Wherever this point is, we can treat it as the fundamental level with respect to the hierarchy of mechanisms.

Ontologically, the crucial point to observe is that mechanisms are not universals but particulars. They are structured collections of parts which occupy a certain region of space and which interact over a certain definite period of time. We characterize these mechanisms by generalizations because very often a mechanism's behavior is repeatable. My body is a mechanism that on repeated occasions interacts with coffee, and coffee repeatedly and reliably has an effect on my behavior. Not only does the very same mechanism exhibit repeatable behavior, but particular mechanisms may be instances of types with consistent behavior. My metabolic mechanisms are broadly similar to those of other human beings and, as a consequence, there will be generalizations (say about the ingestion of caffeine) that hold true of my body and the bodies of many others. But these generalizations are true in virtue of the fact that these mechanisms can and do operate in particular ways on particular occasions, rather than conversely. This understanding of these generalizations also accounts for why they always involve approximations and are only true *ceteris paribus*.

This explanation of the role of causal mechanisms is available so long as the generalizations are mechanically explicable, but here we come to what may seem the key metaphysical issue. If mechanisms are truly going to explain how one event produces another, all of the interactions between parts, at all levels in the hierarchy of mechanisms, will need to be genuinely causally productive. If it were to turn out that these interactions at the fundamental level were not truly interactions, then none of the putative causal relations mediated by mechanisms would be genuine (cf. Psillos 2004, Craver 2007).

We are now concerned with an interaction between two (or more) parts at the bottom of a mechanistic hierarchy. These parts interact (by hypothesis) in accordance with a change-relating generalization or law. But how are we to understand the relationship between the generalization and this interaction? There seem to be three main metaphysical possibilities:

- 1) **Humean Lawlessness** – The interaction nothing more than an instance of a pattern that is described by a generalization.
- 2) **Nomological Determination** – The interaction is governed by the generalization (law).
- 3) **Singular Determination** – The interaction is a singular case of causal determination and any generalizations describing interactions are true in virtue of there being a general pattern of such singular instances.

The first view is the position that fits most naturally with the Mill-Ramsey-Lewis (MRL) view of laws.¹¹ On this view, laws are statements that provide the best balance of simplicity and strength in characterizing events within the world. If there are (as there appear to be in our world) a relatively small number of kinds of fundamental-level parts and some relatively simple generalizations describing how these parts behave in relation to each other, then these generalizations would be obvious candidates for MRL laws. This small set of laws, together with a much larger volume of information about how these parts are organized into hierarchies of mechanisms, will provide a simple and powerful description of the pattern of events in the world. I borrow the term “Humean Lawlessness” from Stephen Mumford (2004), who argues persuasively, that the MRL laws aren’t truly laws, because they supervene on particulars of the actual world. They describe a pattern, but they do not create or explain the pattern. Such a view is anti-realist with respect both to laws and to causes.

The second view – nomological determination – holds that holds fundamental interactions are governed by laws. This view requires some form of nomological realism about fundamental laws, such as the Armstrong-Dretske-Tooley (ADT) view. A law on this view is some third metaphysically real entity, apart from the particular events, which makes it true that one event produces the second. Causal relations are real but subordinate to nomological relations.

The third view, singular determination, holds that there are genuine interactions between parts at the bottom of the mechanistic hierarchy, but that these parts are not governed by laws. In calling these interactions genuine, I am suggesting that the relationship is a modal one. We can express the modality of the relationship counterfactually: When a change in *a* produces a change in *b*, it follows (with the usual caveats about overdetermination, etc.) that if *a* had not changed, *b* would not have changed. But the counterfactual locution should be understood not as a claim about non-actual worlds, but a claim about the determining power of *a* in this world. The singular determination view is the view that is consistent with Anscombe’s and Ducasse’s arguments for singular determination, in the sense that it supports the basic idea that whether a particular event (or change in an object) causes another is at bottom a fact about

¹¹ There has been a good deal of recent literature on alternate interpretations of laws of nature, including the MRL and ADT views. Two helpful introductions are (Psillos 2003) and (Carroll 2008).

the relationship between these two events and is ultimately independent of any facts about other events.

Because the mechanistic approach to causation requires that there be genuine causal connections between parts all of the way down the mechanistic hierarchy, it appears that the question of whether or not the mechanistic approach is a genuinely singularist one would essentially depend upon which of these metaphysical options is correct. If Humean lawlessness is operative at the fundamental level, then mechanisms are simply parts of the pattern of events in the world and they cannot imbue relations between those events with any genuine sort of causal necessity. Alternatively, if fundamental interactions are *governed* by fundamental laws, then the truth of all claims about productive relations between entities at any level in a mechanistic hierarchy will depend upon these laws. Finally, if causal relationships between events at the fundamental level involve singular determination, then so too will events at all levels of the mechanistic hierarchy.

Each of these metaphysical possibilities is genuine in the sense that each of them is consistent with the pattern of objects and events that both science and common experience reveal in the world. If this is the case, which of them should we accept? One approach would be to reject this question as meaningless on the grounds that the options are not empirically distinguishable. While I am not entirely unsympathetic to this sort of anti-metaphysical response, I do think there are arguments that may lend credence to one or more of these metaphysical positions. I cannot hope to survey the many arguments that have been offered in support or criticism of these positions. I can only offer here some explanation of why I think the metaphysics of singularist determination fits naturally with the mechanistic approach to causation that I have argued for.

In the first place, it is difficult to reconcile our intuitions about manipulation, which are central to the mechanistic view, with the Humean view. On the Humean view there is no such thing as genuine modal relationships (or necessary connections as Hume would say) between events. Moreover, singular counterfactual claims are not really claims about what would have happened in a single case. Causal and counterfactual expressions are elliptical ways of talking about complicated patterns in the experience of the actual world. Because of this, manipulations or interventions are not modally effective ways to change the world; they are simply part of the pattern of the actual world. Manipulation, like all other forms of causing, is

shown to be a fiction of the human mind. What patterns or regularities there are in the world just are. There are no explanatory principles to account for them. I concur with Mumford's summary judgment of this sort of metaphysical view – "irrefutable, but neither compelling, appealing nor intuitive" (Mumford 2004: 33). What I think makes it unappealing and non-intuitive is that it is inconsistent with the belief that we manipulate things and cause things in the single case.

The main argument in favor of the nomological determination view is simply that it is implausible to suppose that the order and regularity that we find in nature would exist without laws.¹² Suppose, for the sake of illustration, that patterns of gravitational attractions between bodies are correctly described by the generalizations we call Newton's laws of motion and the law of universal gravitation, and furthermore that these generalizations are not mechanically explicable. Statements of these laws collectively entail certain claims about regular patterns of behavior in the objects – for instance, that satellites will travel around planets in elliptical orbits. It seems quite reasonable, the defender of nomological determination would contend, that these objects behave as they do because they are governed by these general laws. If these laws do not govern their instances, there would be no explanation for the existence of this regularity in nature.

In response to this argument, the singularist must contend that in each particular interaction, the change in one entity produces a change in the other entity, and the fact of this productive relationship does not depend in any way on a general relationship between properties or instances. In the gravitational case, each body acts in each instant on the other body, producing accelerations which, over time, produce elliptical orbits. We do live in a world in which fundamental interactions fit within general patterns, but from this it does not follow that it is in virtue of falling under those patterns that the productive relationship holds. It is possible that we could live in a higgledy-piggledy world in which causes determined effects but in which these determination relations were not regular.¹³

¹² This is what Mumford (2004: ch. 5) refers to as the nomological argument.

¹³ Not all will agree. Heathcote and Armstrong (1991: 63-73), from whom I borrow the term "higgledy-piggledy world", conclude that this situation, while conceivable, is not possible, and that there is an a posteriori necessary connection between laws and causes.

One potential source of confusion has to do with the relevance of properties to causation. A proponent of nomological determination would point out that causal relations appear to hold in virtue of properties of the related events. For instance, the ringing of the alarm clock caused me to wake in virtue of its loudness. Thus, if one takes laws to be relations between properties, it may seem natural to infer that causes depend upon laws. The singularist, however, can argue that on a particular occasion a causal relation between events may hold in virtue of certain properties of those events but not hold to a nomological theory of causation. The singularist does not deny the importance of properties in characterizing causal relations, but insists that it is not essential to causal relations that the relationship between cause and effect be the same on different occasions.

This argument is moreover bolstered by the mechanistic view on the nature of properties and their relations to objects. I have argued (Glennan forthcoming) that much of the literature on causation and laws suffers from a property bias – a tendency to think of causal relations as relations between properties without recognizing how properties themselves depend upon particulars. Consider for example a property of butter – whether it is in a solid or liquid state. Butter is not a basic substance, but a combination of a number of different types of fats – saturated and unsaturated. Whether a fat is saturated or not is in turn dependent upon the molecular structure of the fat. When we say that heating the butter will cause the butter to melt, this is, on the face of it, a change-relating generalization involving properties – changing the temperature of the butter will change the solid/liquid state of the butter. But the properties themselves are not basic facts about the substance butter, but depend instead upon the particular structure of the butter – the molecules that make it up and their arrangement, as well as the arrangement and bonding of sub molecular structures within these fat molecules. One consequence of this is that there is no such property as *the* melting point of butter. Different samples of butter will have different kinds and proportions of fat molecules which will have different consequences for their interaction with heat.

This point about properties applies only to higher-level properties. If an object has fundamental properties like mass and charge whose presence or causal role cannot be explained by reference to the organization and interaction of parts of that individual, then we cannot show how those properties depend upon particulars. For this reason, the observation about property bias is not decisive at the fundamental level. In favor of the singularist

interpretation of fundamental interactions we can only say that it has the advantage of providing a consistent picture of the role of properties and laws in characterizing causal relations.

Where does this all leave us with regard to our original question about the relationship between singular and general causal relations? If the argument has succeeded it has shown that a mechanistic approach to causation is consistent with a singularist causal metaphysics. This is so even though causal generalizations are part and parcel of the apparatus we use to describe and manipulate mechanisms and to formulate causal explanations. Moreover, the singularist picture is the simplest one for the defender of a mechanistic approach to causation, because it fits most naturally with the view that the causal mechanisms which are the truth-makers for causal claims operate at particular locations in space and time. One of the virtues of the mechanistic approach to causation is that it at once fits nicely with a singularist metaphysics and explains the centrality of causal generalization to our epistemic and explanatory practices.

References

- Anscombe, G. E. M. [1971] 1993. Causality and determination. In *Causation.*, eds. Ernest Sosa, Michael Tooley, 88-104. Oxford: Oxford University Press.
- Bogen, Jim. 2004. Analysing causality: The opposite of counterfactual is factual. *International Studies in the Philosophy of Science* 18, (1): 3-26.
- Carroll, John W. 2008. Laws of nature. In *Stanford encyclopedia of philosophy.*, ed. Edward N. Zalta. Fall 2008 Edition ed.
- Craver, Carl F. 2007. *Explaining the brain: What a science of the mind-brain could be.* New York: Oxford.
- Davidson, Donald. [1967] 1993. Causal relations. In *Causation.*, eds. Ernest Sosa, Michael Tooley, 88-104. Oxford: Oxford University Press.
- Dowe, Phil. 2010. Causal process theories. In *Oxford handbook of causation.*, eds. Helen Beebe, Christopher Hitchcock and Peter Menzies. New York: Oxford University Press.
- . 2008. Causal processes. In *The stanford encyclopedia of philosophy (fall 2008 edition).*, ed. Edward Zalta. URL = <<http://plato.stanford.edu/archives/fall2008/entries/causation-process/>> .

- . 2000. *Physical causation*. Cambridge studies in probability, induction, and decision theory. Cambridge ; New York: Cambridge University Press.
- Ducasse, Curt J. [1926] 1993. On the nature and observability of the causal relationship. In *Causation.*, eds. Ernest Sosa, Michael Tooley. Oxford: Oxford University Press.
- Eells, Ellery. 1991. *Probabilistic causality*. Cambridge, MA: Cambridge University Press.
- Fodor, Jerry. 1974. Special sciences (or: The disunity of science as a working hypothesis). *Synthese* 28, (2): 97-115.
- Glennan, Stuart S. under review. Ephemeral mechanisms and historical explanation.
- . forthcoming. Mechanisms, causes and the layered model of the world. *Philosophy and Phenomenological Research*.
- . 2010. Mechanisms. In *Oxford handbook of causation.*, eds. Helen Beebe, Christopher Hitchcock and Peter Menzies, 315-325. New York: Oxford University Press.
- . 2002. Rethinking mechanistic explanation. *Philosophy of Science* 69, (3 Supplement): S342-53.
- . 1997. Capacities, universality, and singularity. *Philosophy of Science* 64, (4) (D): 605-26.
- . 1996. Mechanisms and the nature of causation. *Erkenntnis* 44, (1) (Ja): 49-71.
- Godfrey-Smith, Peter. 2010. Causal pluralism. In *Oxford handbook of causation.*, eds. Helen Beebe, Christopher Hitchcock and Peter Menzies. New York: Oxford University Press.
- Hall, Ned. 2004. Two concepts of causation. In *Causation and counterfactuals.*, eds. John Collins, Ned Hall and L. A. Paul, 225-276. Cambridge, MA: Bradford Book/MIT Press.
- Heathcote, Adrian, and D. M. Armstrong. 1991. Causes and laws. *Noûs* (03/01): 63-73.
- Hitchcock, Christopher. 2007. How to be a causal pluralist. In *Thinking about causes: From greek philosophy to modern physics.*, eds. Gereon Wotter, Peter Machamer, 200-221. Pittsburgh: University of Pittsburgh Press.
- Hitchcock, Christopher Read. 1995. Discussion: Salmon on explanatory relevance. *Philosophy of Science* 62, (2) (June): 304-20.
- . 1995. The mishap at Reichenbach fall: Singular vs general causation. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 78, (3) (June): 257-91.
- Kitcher, Philip. 1984. 1953 and all that: A tale of two sciences. *Philosophical Review* 93, (JL): 335-74.

- Lewis, David. 1973. Causation. *Journal of Philosophy* 70, : 556-567.
- Machamer, Peter. 2004. Activities and causation: The metaphysics and epistemology of mechanisms. *International Studies in the Philosophy of Science* 18, (1): 27-39.
- Machamer, Peter, Lindley Darden, and Carl F. Craver. 2000. Thinking about mechanisms. *Philosophy of Science* 67, (1): 1-25.
- Mumford, Stephen. 2004. *Laws in nature*. New York: Routledge.
- Pearl, Judea. 2000. *Causality : Models, reasoning, and inference*. Cambridge, U.K. ; New York: Cambridge University Press.
- Psillos, Stathis. 2004. A glimpse of the secret connexion: Harmonizing mechanisms with counterfactuals. *Perspectives on Science* 12, (3): 288-319.
- . 2003. *Causation and explanation*. Montreal: McGill-Queen's University Press.
- Russo, Federica, and Jon Williamson. 2009. *Generic versus single-case causality: The case of autopsy*.
- Salmon, Wesley C. 1997. Causality and explanation: A reply to two critiques. *Philosophy of Science* 64, (3) (S): 461-77.
- . 1994. Causality without counterfactuals. *Philosophy of Science* 61, (2) (Je): 297-312.
- . 1984. *Scientific explanation and the causal structure of the world*. Princeton, N.J.: Princeton University Press.
- . 1980. Probabilistic causality. *Pacific Philosophical Quarterly* 61, .
- Sober, Elliott. 1984. Two concepts of cause. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1984, (, Volume Two: Symposia and Invited Papers): 405-24.
- Spirtes, Peter, Clark N. Glymour, and Richard Scheines. 2000. *Causation, prediction, and search*. Lecture notes in statistics. 2nd ed. Vol. 81. Boston: MIT Press.
- Woodward, James. 2004. Counterfactuals and causal explanation. *International Studies in the Philosophy of Science* 18, (1): 41.
- Woodward, James. 2003. *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.
- . 2002. What is a mechanism? A counterfactual account. *Philosophy of Science* 69, (3 Supplement): S366-77.

