Causal Processes and Interactions: What Are They and What Are They Good For?

Christopher Hitchcock†‡

Concerning any object of philosophical analysis, we can ask several questions, including the two posed in the title of this paper. Despite difficulties in formulating a precise criterion to distinguish causal processes from pseudoprocesses, and causal interactions from mere spatiotemporal intersections, I argue that Salmon answered the first of these questions with extraordinary clarity. The second question, by contrast, has received very little attention. I will present two problems: in the first, it seems that Salmon has provided exactly the conceptual resources needed to solve the problem; in the second, it is difficult to see how causal processes and interactions may be used to shed any light. In general, the way to carry Salmon's program forward will be to demonstrate that these resources can be made to do real philosophical work.

For the last quarter century of his life, Wesley Salmon strove to develop an adequate theory of causal processes and interactions. He was particularly concerned to distinguish these from certain impostors: pseudoprocesses and 'mere intersections', respectively. For brevity, I will often speak only of the distinction between causal and pseudoprocesses, even when my comments pertain to both distinctions. In fact, I will do so now. We may ask the following four questions about the former distinction:

†To contact the author, please write to: Division of Humanities and Social Sciences, 101-40, California Institute of Technology, Pasadena CA 91125; e-mail: cricky@caltech.edu.

‡I would like to thank all of those in attendance at the 2002 PSA meeting whose comments—both in the session, and in the bar afterward—helped me to refine the ideas in this paper. I am especially indebted to Peter Machamer for a spirited debate about birth control pills. I also wish to thank my cosymposiasts Phil Dowe and Larry Sklar; Adolf Grünbaum for sharing his personal recollections of Wes Salmon; Paul Humphreys for organizing the symposium; and Merrilee Salmon for giving her blessing to the event. Finally, I would like to thank Wes himself for years of guidance, inspiration, and just plain great philosophy.

Philosophy of Science, 71 (December 2004) pp. 932–941. 0031-8248/2004/7105-0025$10.00

1. Is there a genuine distinction between causal and pseudoprocesses?
2. Are we able to judge when a process is causal, and when it is pseudo?
3. How, exactly, is this distinction to be characterized?
4. To what philosophical use can this distinction be put? Or, to borrow the colorful terminology of Putnam (1962): Does it cut any philosophic ice, bake any philosophic bread, or wash any philosophic windows?

The literature—works by Salmon as well as his critics—has focused almost exclusively on question 3. In fact, Salmon offered a number of different answers to question 3, with several revisions motivated by criticisms. Early versions of his mark-transmission theory appear in several articles published in the late '70s and early '80s (Salmon 1977, 1978, 1982). In response to criticisms from Bas van Fraassen (1980, 123), Salmon clarified the notion of a causal interaction by drawing a distinction between conjunctive and interactive forks. In response to a counterexample from Nancy Cartwright (described in Salmon 1984, 148), Salmon incorporated certain counterfactuals into his characterization of causal processes and interactions. Further criticisms, especially by Philip Kitcher (1989) and Phil Dowe (1992), led him to believe that this involvement with counterfactuals was a serious drawback for his mark-transmission theory. Inspired by Dowe (1992), he reformulated his distinctions in terms of conserved quantities instead.

I do not wish to deny that question 3 is important and interesting. It is certainly the business of the philosopher to provide precise characterizations of rough and ready concepts that others simply take for granted. But there is an opportunity cost to expending one's energy on the formulation and criticism of conceptual analyses, and in this instance, the focus on question 3 has drawn intellectual resources away from the far more important question 4. Philosophers sometimes adopt the conceit that a concept should not be employed in philosophical projects until it has been given a clear and precise definition. I would maintain, by contrast, that an affirmative answer to questions 1 and 2 would suffice for skipping question 3 and proceeding directly to question 4.

Salmon has persuaded me—and he ought to have persuaded *you*, dear reader—that the answer to our first two questions is a resounding 'yes'. In his many writings on the topic, particularly *Scientific Explanation and the Causal Structure of the World* (Salmon 1984), Salmon offered not only his attempts to explicitly characterize the distinction between causal and pseudoprocesses, he also offered a great deal of extraordinarily clear yet informal discussion. For example, he offered a great many illustrations—he taught us that baseballs, cars, arrows, photons, sound waves, and radio signals are all causal processes, while shadows and spots of light on walls

are pseudoprocesses. He taught us that collisions between baseballs or cars are causal interactions, while the crossings of shadows or light beams are mere intersections. Moreover, Salmon showed us how the Special Theory of Relativity *presupposes* the distinction between causal and pseudoprocesses. The first signal principle tells us that 'nothing' can travel faster than the speed of light, or more cautiously, that 'nothing' can accelerate or decelerate across that velocity barrier. Strictly speaking, this principle is false; many things can accelerate across the velocity of light. Suppose that a small but powerful laser located at the center of an enormous 'super-Astrodome' casts a spot of light on the wall of the stadium. Then by rotating the laser with sufficient (but subluminal) velocity we can make the spot on the wall travel faster than light. (Salmon 1984, 141–142) The moral is that we must restrict the scope of the universal quantifier in the first signal principle: it applies only to causal processes, and not to pseudoprocesses. This suggests that the distinction between the two has an objective basis in the physical word.

If one is offering a *stipulative* definition of some concept, then the definition must be crystal clear. It must be possible to determine which items fall under the concept through the mechanical application of the definition. By contrast, if one wishes to point to a distinction that has an objective basis in the world, it suffices to convey enough information through examples and informal discussion to allow competent speakers to classify entities under the appropriate categories, even if that classification relies explicitly or implicitly on background knowledge and not merely on the mechanical application of some definition. Salmon has succeeded admirably in directing our attention to some important physical distinctions.

If I were to present a number of processes, there would be near unanimous agreement among readers of this article about which were causal and which were pseudo. Moreover, where there was disagreement, it would result from ignorance of the underlying physics, not from any conceptual unclarity. (For example, I am ashamed to admit that I *still* get confused over phase velocities and group velocities—I would have a fifty-fifty chance of getting that one wrong.) Our judgment about the status of a particular process is defeasible, of course, but no more so than our judgment about whether, e.g., a given particle is positively or negatively charged.

The question that we really need to be asking is this one: *Given* that there is a genuine distinction between causal processes and pseudoprocesses, what can we do with it? In particular, what sorts of philosophical problems does it help us to solve?

Here is one type of problem where I think that the concept of a causal process helps us enormously. In his paper "On the Notion of Cause"

(1913), Bertrand Russell argued that there is no such thing as causality. While this paper is no doubt best known for Russell's irreverent poke at the British monarchy, it also posed some interesting Zenoesque paradoxes concerning causation. These paradoxes are addressed to the following definition, taken from Baldwin's *Dictionary*:

> Cause and effect . . . are correlative terms denoting any two distinguishable things, phases, or aspects of reality, which are so related that whenever the first ceases to exist the second comes into existence immediately after. . . .

This definition involves two different components: temporal contiguity and invariable succession. Russell's critique of this definition runs along the following lines: If the cause and effect—call them *C* and *E*—are temporally contiguous and distinct, then at least one of them must be temporally extended. For if they are both instantaneous, point-like in time, then they cannot be contiguous. (Just as two numbers on the real number line must be separated by some finite distance.) Assume, for the sake of argument, that *C* is extended. (Imagine, if you like, that *C* corresponds to the half-closed set (0, 1], while *E* is at least partially open, with 1 as its greatest lower bound.) But now, by contiguity, only the second half of *C* is really efficacious—for the first half of *C* is at a temporal remove from *E*. Repeat this argument for the final three-quarters, seven-eighths, and so on, and we have an obvious regress. The upshot is that at most one instantaneous temporal slice of *C* can be efficacious, contradicting our assumption that the cause is temporally extended. (An analogous reductio can be constructed if we assume instead that *E* is temporally extended.)

Russell's puzzle arises because the cause and effect are assumed to be discrete entities. Within this framework, the only way to capture the idea that causation is *local* (there is no causation at a spatiotemporal distance) is to require that causes and effects be contiguous with one another. One could have a series of events—a causal chain—linking *C* and *E*. But then the same problem will arise in connection with successive links in the chain. The notion of a causal process seems to be tailor-made to capture the spatiotemporal locality of causal influence while avoiding the paradox that results from assuming discrete events to be contiguous.

Salmon was certainly interested in Zeno's paradoxes; he did, after all, edit a book on the topic (Salmon 1970). His mark-transmission theory of causal processes was heavily influenced by Russell's 'At-At' solution to Zeno's paradox of the arrow (see Salmon 1977; 1984, 147–157). There are even a number of places where Salmon hints that the notion of a causal process can be used to avoid problems that arise from the assumption that causes and effects are discrete entities (e.g., 1984, 156–157,

182–183). So far as I am aware, however, Salmon never explicitly addressed Russell's Zenoesque puzzle. I harbor little doubt, however, that Salmon has provided exactly the conceptual resources needed to solve this puzzle. The exact formulation might take some care: perhaps some clever graduate student could take up the project and get a nice little publication to help her on the job market.

In contrast to this puzzle, let us consider the problem of explanatory irrelevance, as illustrated in Salmon's famous example:

> John Jones avoided becoming pregnant during the past year, for he has taken his wife's birth control pills regularly, and every man who regularly takes birth control pills avoids pregnancy. (Salmon 1971, 34.)

Salmon presents this as a counterexample to Hempel's Deductive-Nomological model of scientific explanation (Hempel 1965). John Jones did indeed consume birth control pills (or so we may suppose), it is indeed a law that no man who consumes birth control pills will become pregnant, and Jones's avoidance of pregnancy does follow deductively from the preceding premises. Despite all of this, we do not have a successful explanation of Jones's barrenness. The problem is that the putative explanation provides information that is *irrelevant*: Jones's consumption of birth control pills has nothing to do with his failure to conceive.

Salmon's example wants for nothing by way of memorability. It does, however, contain one small red herring. Birth control pills *do* often prevent pregnancy in women, but this is in no way a crucial feature of the example. We would have an equally persuasive counterexample if John Jones had consumed aspirin or broccoli instead of birth control pills. For the sake of tradition, I will continue to use Salmon's original example, but it is important to keep this point in the back of our minds.

This famous example first appeared in Salmon's paper "Statistical Explanation" (1971), which introduced his statistical-relevance model of explanation. That theory did indeed offer an effective resolution of this example: the class of human men is objectively homogeneous with respect to pregnancy; subdividing the class of men into those that consume birth control pills and those that do not will not effect a statistically relevant partition. It is clear that, in this paper, Salmon took statistical relevance relations—at least statistical relevance relations within otherwise homogeneous reference classes—to be intrinsically explanatory. They were explanatory, Salmon believed, because they were closely connected with causal relevance relations. In fact, he explicitly expressed the hope that it would eventually be possible to analyze causation in terms of probability. Throughout the 1970s, however, Salmon became more pessimistic about the prospects for developing an adequate probabilistic theory of

causation. This process culminated in 1980 with the publication of an influential critique of probabilistic theories of causation (Salmon 1980). Statistical relevance relations, he concluded, could provide evidence for the existence of causal relations, but they were not constitutive of causal relations. From the earlier statistical-relevance theory, Salmon maintained the idea that explanatory relevance is causal relevance, while dropping the idea that this latter notion could be articulated in terms of probabilities. It was during this same period of time that Salmon began working on his theory of causal processes and interactions, in the hopes that it would succeed where probabilistic theories had failed. The problem of explanatory irrelevance thus played an important, albeit indirect, role in leading Salmon to his mature theory of causation.

A natural question, then, is how the conceptual resources of Salmon's new theory of causation can help us to solve the problem of explanatory relevance. The answer to this question is not at all obvious. Birth control pills are causal processes, after all: they are capable of transmitting marks and conserved quantities. When John Jones swallows a birth control pill, that is a genuine causal interaction. (More precisely, it is what Salmon calls a λ-type interaction, where two incoming processes merge into a single outgoing process.) So the problem with the putative explanation is *not* that it mistakes a pseudoprocess for a causal process, or a mere intersection for a causal interaction. So just how are the concepts of causal process and interaction to be of help in explicating what is wrong with this explanation?

In response to this challenge, defenders of process theories of causation (and even some of their critics) have offered a number of replies (either singly or jointly), which I will consider.

> **Reply 1.** We know what processes take place in female bodies when birth control pills *do* prevent pregnancy, and we know that these processes can not take place in a male's body.

This response results from following the scent of the red herring described earlier. This response cannot be correct: if we replace Salmon's example by one in which John Jones takes aspirin or eats broccoli, the first reply is simply unavailable. But surely the basic reason for explanatory irrelevance is the same in each of these cases.

A second, related, reply avoids this problem:

> **Reply 2.** We know what processes and interactions take place within male bodies, and we know that those processes are incompatible with pregnancy, so we know that the addition of birth control pills must be irrelevant.

There are at least two problems with this reply. The first is that it is not

at all clear what the appeal to processes and interactions adds that is not already present in the response provided within the framework of Salmon's statistical-relevance model. Men cannot get pregnant, period; therefore, the consumption of birth control pills cannot make a difference to whether or not a man becomes pregnant; therefore, the consumption of birth control pills is not explanatorily relevant to John Jones's pregnancy. This line of response is not only available within Salmon's statistical-relevance theory, but also within a counterfactual theory of explanation (Lewis 1986, Woodward and Hitchcock 2003, Woodward 2003), as well as *any* account that defines explanatory relevance in terms of necessary or even probabilistically relevant conditions. Is anything really clarified by adding that a man's inability to become pregnant is due to various internal processes?

The second problem is more subtle and yet more profound. In summarizing his famous covering-law model of explanation, Carl Hempel wrote:

> scientific explanation . . . seeks to provide a systematic understanding of empirical phenomena by showing that they fit into a nomic nexus. (Hempel 1965, 488)

Salmon summarized his own view by modifying this passage:

> my suggestion for modification would be to substitute the words '*how* they fit into a *causal* nexus' for '*that* they fit into a *nomic* nexus'. (Salmon 1984, 19)

A causal nexus, for Salmon, is a network of causal processes and interactions. More specifically:

> If we want to show why *E* occurred, we fill in the causally relevant processes and interactions that occupy the past light cone of *E*. (Salmon 1984, 19)

Of course, what's at issue here is how to determine which processes and interactions *are* causally relevant, but let us bracket this for just a moment. The point I wish to stress is that explanation, for Salmon, involves the exhibition of *actual* causal processes and interactions in the *causal history* of the event or phenomenon to be explained. In our example, the event to be explained is John Jones's failure to become pregnant during a specific one-year period. Now how on earth can the causal processes and intersections that take place within *other men* have any bearing on this? For one thing it does not seem to be crucial to the example that such processes and interactions have even occurred in Jones's past light cone: would the counterexample not have worked equally well if John Jones had been replaced by Adam? But even if such processes and interactions had taken place, these seem to be exactly the sorts of processes and interactions that

should be deemed *irrelevant* to Jones's present state. The goings-on beneath the skin of other males surely did not interfere with conception in Jones's body!

If the causal processes and interactions that take place within the bodies of other men have any bearing, it is at best indirect. For example, they may provide us with information about the statistical relevance of birth control pill consumption for pregnancy in men. Salmon makes it very clear in his later work that statistical-relevance relations are not *constitutive* of causal (and hence explanatory) relations, but can only *provide evidence* for the genuinely causal relations consisting of processes and interactions (see e.g. Salmon 1984, 34n10, 192, 265.) Alternatively, we may take the processes that occur within other males to provide evidence about what *would* have happened in Jones's body had he not consumed his wife's pills. In *Scientific Explanation and the Causal Structure of the World*, Salmon grudgingly allowed that counterfactuals play a role in determining whether a process is causal or not (1984, 148–150); but even so, it is the processes themselves that figure in genuine explanations, and not the attendant counterfactuals (see Hitchock 1995, sect. 4, for more discussion). In the conserved quantity theory proposed by Salmon (1994), counterfactuals are quite deliberately expunged from the account altogether.

The upshot, I think, is that the processes and interactions that take place within the bodies of other males can only have a bearing on the example insofar as they provide evidence about the processes and interactions that actually occur in Jones's case.

Have we not just suggested another reply to our challenge? Here it is:

> **Reply 3.** The irrelevance of the birth control pills for John Jones's failure to conceive is manifest in the network of microscopic processes and interactions that take place within Jones's body.

This reply is unsatisfactory for at least three reasons. First, we are able to recognize the irrelevance of pill consumption without any detailed knowledge of the processes that take place when a male consumes birth control pills. (At least I lack such detailed knowledge; but I suspect that you do too, dear reader.)

Second, it seems to me that the original problem is likely to reappear at the microlevel. When we consider the explanandum, Jones's failure to become pregnant during the past year, it is a little difficult to locate this failure in space and time. If it cannot be assigned a spatiotemporal location, then it will not be possible to find *any* Salmon-style explanation of it, for it will have no past light cone in which to search for relevant processes and interactions. So let us say that Jones's nonpregnant state does have a location: it occurs within some region of Jones's body (or

perhaps his entire body) during the particular span of time in question. Now, when Jones consumes birth control pills, those pills enter his stomach and are dissolved by hydrochloric acid. The chemical contents of the pills will be transported throughout his body in his bloodstream, and hence into the region where his nonpregnancy is supposed to occur. These are genuinely causal processes and interactions, and they connect his original interaction with the pills (his swallowing them) with (various later parts of) his nonpregnant state. So why do these processes and interactions not count as 'relevant' to Jones's failure to become pregnant? Our original problem has simply been hidden under Jones's skin. (See also Hitchcock 1996, sect. 4, for related discussion.)

Finally, there is a more general problem. That Jones's consumption of birth control pills is irrelevant to his failure to become pregnant surely has *something* to do with the fact that he would have avoided pregnancy even if he had not taken the pills. We may hope to eliminate the modality by analyzing it in terms of regularities, correlations, or what have you; but the surface modality clearly seems to be central. This idea is captured neatly in Salmon's statistical-relevance model of explanation. Yet a specification of the *actual* processes and interactions that took place tells us nothing about what *would have* happened if those processes had not occurred. At best it can give us counterfactual information of the following sort: There is some conserved quantity Q, such that if the amount of Q within a certain spatiotemporal region had been q, then it would also have been q within a certain nearby spatiotemporal region as well. Similar comments apply to the mark-transmission version of Salmon's theory (see Hitchcock 1995, sect. 4). These counterfactuals, however, assume that the actual causal processes (or near counterparts of them) are present within the hypothetical scenario described. We are simply unable to say, using only the language of causal processes and interactions, that Jones would not have become pregnant even if he had not taken the pills.

The problem of explanatory relevance is clearly one that Salmon took to be of central importance. I have raised the question of whether Salmon's concepts of causal process and interaction can be put to use to shed light on this problem. I am highly skeptical, but my goal in this essay is not really to argue that it *cannot* be done. For example, Phil Dowe, in chapter 7 of his *Physical Causation* (2000; see also his 2004 in this volume), offers a proposal that that may bear on this issue, and I have taken no pains to refute that particular proposal here. My goal, rather, is to illustrate a certain *type* of question that has received very little attention in the vast literature inspired by Salmon's work: just what sorts of philosophical problems *can* be successfully addressed, once one has taken Salmon's conceptual apparatus on board? The most fitting philosophical tribute to Salmon's memory is to make genuine progress on those issues that he

cared about. With respect to causal processes and interactions, the way forward will not involve attempts to dot the *i*'s and cross the *t*'s in the definitions of these concepts. Rather, we need to show how these concepts can be made to do real philosophical work.

REFERENCES

Dowe, Phil (1992), "Wesley Salmon's Process Causality and the Conserved Quantity Theory", *Philosophy of Science* 59: 195–216.

——— (2000), *Physical Causation*. Cambridge: Cambridge University Press.

——— (2004) "Causation and Misconnections", *Philosophy of Science* 71 (Proceedings): 926–931.

Hempel, Carl G. (1965), "Aspects of Scientific Explanation", in *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press, 331–496.

Hitchcock, Christopher R. (1995), "Discussion: Salmon on Explanatory Relevance", *Philosophy of Science* 62: 304–320.

——— (1996), "The Mechanist and the Snail", *Philosophical Studies* 84: 91–105.

Kitcher, Philip (1989), "Explanatory Unification and the Causal Structure of the World", in Philip Kitcher and Wesley C. Salmon (eds.), *Scientific Explanation.* Minneapolis: University of Minnesota Press, 410–505.

Lewis, David K. (1986), "Causal Explanation", in *Philosophical Papers,* vol. 2. Oxford: Oxford University Press, 214–240.

Putnam, Hilary (1962), "The Analytic and the Synthetic", in Herbert Feigl and Grover Maxwell (eds.), *Scientific Explanation, Space, and Time,* Minnesota Studies in the Philosophy of Science, vol. 3. Minneapolis: University of Minnesota Press, 358–397.

Russell, Bertrand (1913), "On the Notion of Cause", *Proceedings of the Aristotelian Society* 13: 1–26.

Salmon, Wesley C. (1971), "Statistical Explanation", in Wesley C. Salmon et al., *Statistical Explanation and Statistical Relevance*. Pittsburgh: University of Pittsburgh Press, 29–87.

——— (1977), "An 'At-At' Theory of Causal Influence", *Philosophy of Science* 44: 215–224.

——— (1978), "Why Ask 'Why?'?" *Proceedings and Addresses of the American Philosophical Association* 51: 683–705.

——— (1980), "Probabilistic Causality", *Pacific Philosophical Quarterly* 61: 50–74.

——— (1982), "Causality: Production and Propagation", in Peter Asquith and Ronald Giere (eds.), *PSA 1980,* vol. 2. East Lansing, MI: Philosophy of Science Association, 49–69.

——— (1984), *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.

——— (1994), "Causality without Counterfactuals", *Philosophy of Science* 61: 297–312.

——— (ed.) (1970), *Zeno's Paradoxes*. Indianapolis: Bobbs-Merrill.

van Fraassen, Bas C. (1980), *The Scientific Image.* Oxford: Clarendon Press.

Woodward, James F. (2003), *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

Woodward, James F., and Christopher R. Hitchcock (2003), "Explanatory Generalizations, Part I: A Counterfactual Account", *Noûs* 37: 1–24.