Causality

An analysis of causation aims to achieve an understanding of what it is for two events to be related as cause and effect – what *must* be the case, metaphysically speaking, for two events to be so related. This kind of inquiry is different from the inquiry by which we come to *know* that two events stand in the causal relation, which is an epistemological project.

There are many kinds of analyses of causation in the literature. And as one of the most important and widespread issues in philosophy, there is little agreement about the correct approach.

Some philosophers say a cause makes something happen, brings something about or makes a difference; and "the difference it makes must be a difference from what would have happened without it," (Lewis 1983: 557). When we say that the financial crisis could have been avoided, we mean in part that there was an event, or were events, without which the crisis would not have happened. These are the causes of the crisis.

A number of issues reappear in the discussions below. As we will see, differentiating between what is causal and what is only apparently causal is an important task for every theory. That is, we often must ask ourselves if we have epistemically arrived at a cause or at an event merely correlated with the cause. And we must ask ourselves, metaphysically speaking, what makes the difference between a cause and an event merely correlated with the cause.

Finally, causation is implicated in philosophical theories in almost every philosophical domain. In ethics, for example, causation might make the difference between killing and letting die. There is the causal theory of reference and the causal theory of knowledge. In this entry, we examine the special role of causation in the mental causation debate in the philosophy of mind, as well as the role of causation in the problem of freewill and determinism.

Table of Contents

1. The Nature of Causality
    a. Causation versus Correlation
    b. Causation versus Causal Explanation
        i. Theories of Scientific Explanation
    c. Causation, Possibility and Necessity

2. Causation in the History of Philosophy
    a. Aristotle's Four Kinds of Causation
    b. Hume
    c. Kant

3. Causal Analyses
    a. Regularity Theories
    b. Mackie's INUS Conditions
    c. Counterfactual Theories
    d. Probabilistic Causation
    e. Process Theories
    f. Manipulationist Accounts

4. Special Issues in Causation
    a. Mental Causation
    b. Free Will and Causal Determinism

1. The Nature of Causality

Consider these two statements:

(a) The short circuit caused the fire.
(b) Short circuits cause fires.

The former is an example of what philosophers call token or singular causation – an individual event causing another individual event. The latter is an example of type or general causation – causation between properties or kinds of events. Many philosophers hold that token and type causation are different kinds of causality. Sober (1985, 1986) and Eells (1991), for example, argue for theories of type and token causation that differ considerably. Cartwright (1989) and Hausman (1998) hold that type causation depends upon the more fundamental token causation. On the other hand, it is possible to hold the view that type-level relations are

fundamental (see Tooley 1987). In this entry, we will be talking about token causation unless otherwise announced.

a. Causation versus Correlation

Often the first word on causation is "Correlation is not causation." It is true that the difference between genuine causal connections and mere correlations is central to an understanding of causation. However, the saying, "Correlation is not causation," is incomplete. It would be more accurate to say, "Correlation is necessary but not sufficient for causation." Two events, c and e, may be constantly correlated and not yet causally related.[1] For instance, they may have a common cause in a third event, b.

Consider this example: sleeping with the lights on is strongly correlated in adults with waking up with cephalalgia, or headache. One might conclude, epistemologically speaking, that sleeping with the lights on caused the headache. However, if we can imagine a plausible common cause then this inference will seem illegitimate. After all, it is possible that excess alcohol consumption caused both the headache and sleeping with the lights on. Making this epistemological error – the "common cause error" – amounts to mistaking an effect for a cause. Now, metaphysically speaking, for the assertion that Event 2 is a cause of Event 3 to be true, there *must not* be another event, Event 1, causing both Event 2 and Event 3. In Figures 1 and 2 the dotted arrow represents an apparent connection, which is in fact no connection at all.

[1] Yet any genuinely causally related events are also constantly correlated (and this is so in virtue of being causally related).

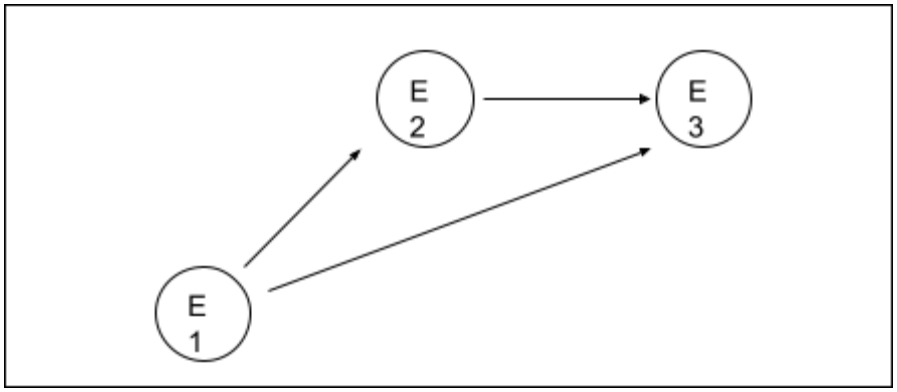

Consider a more concrete example. Nietzsche writes of the strange case of Cornaro, who thought that his meager diet caused longevity. Nietzsche says, "The worthy Italian thought that his diet was the cause of his long life, whereas the precondition for a long life, the extraordinary slowness of his metabolism, the consumption of so little, was the cause of his slender diet,"[2] as well as of his longevity.

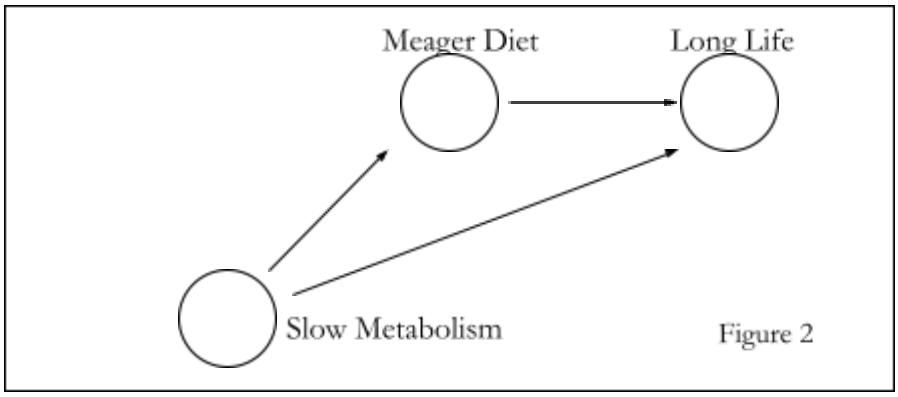

b. Causation versus Causal Explanation

Consider these two token causal claims:

(a) The hurricane caused the building to collapse.
(b) The event reported on page five of the *Times* caused the building to collapse.

[2] For more see Leiter (2002: 156-157).

Assuming (a) and (b) are both true, only (a) is explanatory; (b) is not explanatory, even assuming, as we do, that the hurricane is the very same thing as the event reported on page five of the *Times*. True singular causal statements are called "extensional" because they remain true under substitution of all true descriptions of the relata (the things related by the causal relation). The hurricane caused the building to collapse and so did the event on page five of the *Times* because they are one and the same thing. On the other hand, causal *explanation* can be called "intensional" because there are salience requirements on the proffered descriptions of the causal relata in order for the explanation to be genuinely explanatory.

If we ask for an explanation of why Oedipus gouged out his eyes, it would *not* be explanatory to say, "because Oedipus lusted after the Queen." That description is not relevant or salient. However it *would* be explanatory to say, "because Oedipus lusted after his mother," because when you do that and you're a Greek hero, *it makes sense* to gouge out your eyes. Causal explanations, Davidson (1967) says, are intensional in this sense: they need not preserve explanatory adequacy upon substitution of co-referential terms (differing terms that nevertheless have the same referent). Even though Jocasta, his mother and the Queen, are all the same entity, only under some descriptions does reference to Jocasta explain Oedipus's gouging out his eyes. Nevertheless, it is *true* that he lusted after Queen Jocasta/his mother just as it is true that the event reported on page five of the *Times* caused the building to collapse. Each is true, but not explanatory. Causation and causal explanation are different insofar as it is possible for a sentence to be non-explanatory while still being causal.

Interest-relative, mind-dependent features of a situation are relevant to whether an explanation is more or less explanatory. Lewis's (1979) theory of causal explanation includes the idea that an explanation describes (or gives information about) a part of the causal history of the event to be explained. The relevant part of the history will change relative to the

context of the inquiry. Imagine a car crash. The road engineer will find the improper grading for the arc of the road's curve to be explanatory. The policeman will find the driver's speed to be explanatory. The driver himself, having safely navigated the curve at that speed many times before, may blame the fresh rain on the road.

As Lewis says:

> If someone says that the bald tire was the cause of the crash, another says that the driver's drunkenness was the cause, and still another says that the cause was the bad upbringing which made him so reckless, I do not think any of them disagree with me when I say that the causal history includes all three. They disagree only about which part of the causal history is most salient for the purposes of some particular inquiry, (1979: 215).

On the other hand, it is often considered a desideratum that causation per se be mind-independent. "Causal relationships are features of the world: they are 'out there' in nature. By contrast, explanation is an activity carried out by humans and conceivably by some other animals, having to do with the discovery and provision of information, information based on causal relationships," (Woodward 2003: 23).

i. Theories of Scientific Explanation

Philosophers have disagreed about the extent to which theories of scientific explanation should make fundamental reference to causal concepts. Remarking that "the word 'cause' never occurs" in "advanced sciences such as gravitational astronomy," Russell quipped that "The law of causality… like much that passes muster among philosophers, is a relic of a bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no harm," (Russell 1913). If for the empiricist tradition (of which Russell was a part) causal concepts are not allowed as primitives, then either causal concepts must be reduced to empiricist-pedigreed concepts or they must be avoided by empiricists.

So early on, logical empiricists tried to do without causal concepts by employing the deductive-nomological (DN) model of scientific explanation. More recently, problems with the DN model suggest that theories of scientific explanation *require* causal concepts. DN, as developed by Hempel & Oppenheim (1948) and Hempel (1965), has it that explanations are deductive arguments. The premises of the argument are the explanans (the thing doing the explaining) and the conclusion of the argument is the explanandum (the thing explained). The explanans contains at least one law of nature and it must contain it essentially such that the inference would not go through were it left out. An explanandum is explained if it is possible to *deduce* it from sentences about the antecedent conditions together with the law of nature. To use Hempel's example, the Dust Bowl farmers migrated to California because – and here's the law – "populations tend to migrate to areas offering better living conditions" and California offered better living conditions – there's an initial condition. The explanatory law predicts their migration.[3]

Among many well-known problems,[4] DN has difficulty accounting for what we usually take to be the asymmetry of many scientific explanations. For instance, we normally take there to be an asymmetry pointing from the explanans to the explanandum: we derive the length of a shadow from the height of a flagpole and the angle of the sun (plus a law). The DN model captures this explanation perfectly. But it just as perfectly goes the other way: it also purports to explain the height of the flagpole in terms of the angle of the sun and the

[3] Hempel augmented and made his theory more sophisticated by allowing that historical explanations were often only "explanation sketches" which would account for the poor predictive capabilities of the lessons of history. Except "Winter invasions of Russia always fail," there are few laws in history capable of providing true predictions. Hempel also suggested statistical laws for historical explanations that would predict and explain the frequency of kinds of events.

[4] For a comprehensive discussion of the main problems for DN see Salmon (1989).

length of the shadow (plus the law).[5] But this is a counterintuitive result. DN is allowing too many things to count as explanations. We do not normally explain the height of a flagpole by referencing the length of its shadow. Normally, it goes the other way around.[6] Many philosophers argue that what DN is missing is the asymmetry of causation. Kitcher (1989) diagnoses DN's symmetry to be an unhappy result of the empiricists' Humean desire to do without the "mysterious" concept of causation.[7]

c. Causation, Possibility and Necessity

One way to think of causation is to think that a cause necessitates its effect. That is, given the cause you would necessarily get (it would be impossible not to get) the effect. Causation is, thereby, intimately bound up with possibility and necessity. Some explanation of these terms is required.

Possibilities are things that could have happened, even if they never take place. Something is necessary when it is impossible that it fail to happen. Necessity and possibility are, thereby, like two sides of a coin. Furthermore, there is nomological or natural necessity and natural possibility, and then there is logical necessity and logical possibility. For example, a round square is logically impossible while faster than light travel and 3 cubic mile gold cubes are naturally impossible but logically possible. Insofar as the concept of causation involves necessity, it is natural necessity and not absolute, or logical, necessity. That is to say, the relation of necessity between cause and effect is a relation of necessity only in this world and not in every possible world. That is, things might have been different.

Some historians of philosophy argue that Spinoza and Leibniz collapsed causal relations and logical relations such that the events that have happened could not have failed to have happened or could not have happened otherwise. For necessitarianism of this kind,

---

[5] See Bromberger (1966).

[6] However, see van Frassen (1980: 132-134) for a thought experiment the context of which makes it seem legitimate.

[7] We discuss Hume on causation below.

falsity overlaps with impossibility.[8] Hume's decisive break from necessitarianism has been extremely influential. "All events seem entirely loose and separate," Hume says.

2. Causation in the History of Philosophy

a. Aristotle's Four Kinds of Causation

For Aristotle, it is usually things or substances that have causes, not events. The Greek word, *aition*, which is often translated as "cause," actually means, "the being of a thing." For Aristotle there are as many senses of causation as there are senses of "being" to be explained. A thing's material cause is the matter that constitutes the thing. The formal cause is the thing's essence. The thing's efficient cause is a precipitating antecedent event often mechanistically related to the thing. The final cause is the end, *telos*, or purpose of the thing. So, famously, for Aristotle the oak is the final cause of the acorn. Organic matter is the material cause of the acorn. The formal cause is what it is to be an acorn, while the efficient cause is the tree that dropped the acorn.

According to Dijksterhuis (1961), the New Science of the early modern period ushered in efficient causation as the only legitimate form of scientific explanation. Descartes especially wanted to do away with final causes, a kind of explanation he associated with the Aristotelian schoolmen's moribund philosophy. More broadly, the culture at large was suspicious of occult powers. Molière found it ridiculous to explain why opium makes one sleepy by referring to opium's "dormitive virtue" (*virtus dormitiva*), in other words its power to make one sleepy.

[8] For more on Spinoza's modal metaphysics see Newlands (2007).

b. Hume

Hume's work on causation is perhaps the most influential in all philosophical thought on the topic. Hume writes:

> We may define a cause to be *an object followed by another, and where all the objects, similar to the first, are followed by objects similar to the second.*
>
> Or, in other words, *where, if the first object had not been, the second never had existed.*

Philosophers have found inspiration for regularity theories in the first definition and for counterfactual theories in the second.[9] As above, the logical positivists read Hume to be a regularity theorist, such that relations of cause and effect exist "out in the world," but lack necessity and are nothing more than constant conjunction. Still today philosophers call a position "Humean" if it denies physical necessity, "out in the world," and denies that we have any idea, "in our minds," of necessary connection. For these so-called Humeans the notion of one event *necessarily* entailing the existence of another event is mysterious. Positions of this kind continue to be referred to as Humean positions, even though more and more historians of philosophy now agree that the real, historical Hume would be at considerable variance from these Humeans.

> Off the table now is the idea that [Hume] taught a strict regularity theory: that there is nothing in reality but regular sequence, and that that is accordingly all that causality amounts to, either in our concept of it or in things and events themselves. True, this tendency to speak of regularity theories as 'Humean' persists...[but] this usage just betokens a limited acquaintance with the work of Hume, (Craig 2000: 113).

[9] A standard regularity theory will start by saying: c caused e if and only if c preceded e, and events like c are always followed by events like e. See Section 3a for more detail. A standard counterfactual theory will start by saying: c caused e if and only if c preceded e, and if c had not occurred then e would not have occurred. See Section 3b for more detail.

Hume, in fact, subordinated his theory of causation to his radically empiricist theory of ideas. Perceptions, or what Hume calls "the elements" of his philosophy, came in two kinds, impressions and ideas. Impressions are the more vivid of what is "present to the mind," and they are causally prior to ideas. Ideas are less vivid and causally dependent on impressions. Every idea must be traceable to an original sense impression in order to be legitimate. Hume argued that neither the senses nor reason can establish the idea that one object necessarily entails the existence of another object. According to Hume's picture, the relation between cause and effect has only three perceptible features, (1) contiguity in space and time, (2) temporal priority of the cause, and (3) constant conjunction in the past. And yet there seems to be more to our experience of the causal relation. Indeed, Hume says, "There is a necessary connexion to be taken into consideration," (1.3.2.5-11).

It is part of our nature that regularities of this sort are accompanied by a feeling and it is this feeling from which the idea of necessity is copied.

> After a frequent repetition, I find, that upon the appearance of one of the objects, the mind is determin'd by custom to consider its usual attendant… 'Tis this impression, then, or determination, which affords me the idea of necessity, (1.3.14).

The idea of necessity does not have a proper empiricist pedigree in the theory of ideas: it cannot be traced back to a sense impression. But we nevertheless come to expect like from like; and to expect same cause, same effect. Hume suggests a skeptical outlook with respect to the power of reason to allow us to understand *that thing in reality in virtue of which causes are causal.* Anything beyond constant conjunction is beyond human understanding although we project onto the world our idea of necessity acquired through habit and custom.

c. Kant

For Kant, causation justifies experience rather than the other way around. Kant gives a transcendental argument such that the necessary conditions of possibility of experience include the objective necessary connection between some objects or events.

> The very concept of cause so obviously contains the concept of a necessity of connection with an effect and a strict universality of the rule, that the concept [of cause] would be entirely lost if one pretended to derive it, as Hume did, from a frequent association of that which happens with that which precedes, and [from] a thereby arising custom (thus a merely subjective necessity) of connecting representations, (B5).

Kant argues that if a world of objects is to be represented, then we must be able to distinguish between the subjective and objective character of representations. We need to be able to distinguish, Kant says, between a succession in the objects themselves and the succession occurring in the subject's representations. In order to establish that our representations are about objects, the relation of their succession must be "determined," meaning that it must be necessary and irreversible (A198-9/B243-4). To illustrate this point Kant contrasts two cases, one in which our experience changes while an object remains the same and one in which our experience changes *because* the object itself is changing. The first is the perception of a house in parts and the second is the perception of a boat moving downstream. "The apprehension of the manifold in the appearance of a house that stands before me is successive," Kant says, "but no one will concede" that "the house itself is also successive," (A190/B235). In other words, it is the subject's own actions that determine the order of succession, depending on which part of the house the subject looks at first, second and so forth. "My perceptions could have begun at its rooftop and ended at the ground, but could also have begun below and ended above," (A192/B237). So the order of succession of the representations of the house is best thought of as subjective and not objective and necessary.

By contrast, such reversals of order are not possible in the boat case. "My perception of its position downstream follows my perception of its position upstream," such that "it is impossible that in the apprehension of this appearance that the ship should first be perceived downstream and afterwards upstream," (A192/B237), which makes the order necessary. As Gardner says, "the concept of a necessary and irreversible succession is, Kant says, the concept of a causal relation," (1999: 176). Kant takes himself to have shown that the concept of causation genuinely applies to objects outside us in precisely the way Hume denied.

3. Causal Analyses

a. Regularity and Nomological Theories

A standard regularity theory of causation will start by saying: c caused e if and only if c preceded e, and events like c are always followed by events like e. This naïve definition captures the intuitions behind sayings like "The same cause is always accompanied by the same effect" and Hume's first definition: "an object followed by another, and where all the objects similar to the first are followed by objects similar to the second." In a bit more detail:

c causes e if and only if:
- (1) c precedes e in time
- (2) c and e are contiguous in space
- (3) c is a C-type of event
- (4) e is a E-type of event and
- (5) C-type events are regularly followed by E-type events.

There is a Humean tradition of regularity theories in which the regularities in (5) are merely de facto regularities without any aspect of necessity. Another kind of regularity theory is the nomic subsumption view, which says that when c is a cause of e this is so in virtue of a relation of necessity or a strict law between C-type events and E-type events. So (5) would not be merely de facto regularity but regularity with nomic necessity.

One problem for regularity theories focuses on cases of preempted backups. For example, there is a regularity between the hospital's lights being on and the availability of a backup source of power, the same as there is a regularity between the lights being on and the functioning of the hospital's normal power source. In other words, the de facto regularity theory above cannot tell the difference between mere correlations and genuine causal relations. The nomic subsumption view stipulates the difference but then must account for the distinction between laws of nature and mere accidental generalizations, which is no easy feat.

b. Mackie's INUS Conditions

Above we have sometimes spoken of "*the* cause." But do not most events have multiple factors that bring them about? Once you start thinking of cause*s* – plural – instead of *the* cause, you see that there is a set of conditions that were jointly sufficient but not necessary for the effect. For example, a fire can be started by a short circuit but only in the presence of oxygen. Together those two states amount to a sufficient but not necessary condition for a fire. That is, they are enough to cause a fire, but the fire could be caused in another way, say by striking a match. With respect to the cause of the fire, Mackie (1975) offered a proposal he called INUS conditions:

> In this case, then, the so-called cause is… an insufficient but necessary part of a condition which is itself unnecessary but sufficient for the result, (414).

So the short circuit is itself a necessary part of the set of conditions which were sufficient but not necessary for the fire. Furthermore, the short circuit is, with respect to the set of conditions of which it is a part, insufficient alone for the fire. It needs the rest of the set of conditions. So the cause, the short-circuit, is an *insufficient* but *necessary* part of a set of conditions, the set itself being *unnecessary* but *sufficient* for the result – INUS. Thus, C will be an

INUS condition for E if there is a set of factors ABCD… such that whenever ABCD… occur together, they are invariably followed by E. And it must be the case that when ABD… occur without C, these conditions are not invariably followed by E.

One problem with this account is that it yet may be merely an accidental correlation that whenever ABCD… occur they are followed by E. The INUS account's difficulty in differentiating a mere correlation from a genuine causal relation could be averted by affirming that E follows ABCD… according to a law of nature. However, this again only pushes the problem back a level: now we must solve the recalcitrant problem of differentiating non-causal, accidental regularities from genuine laws of nature.

c. Counterfactual Theories

A simple counterfactual analysis of causation holds that c causes e if, had c not occurred, then e would not have occurred. Although this is overly simplistic, causation is clearly bound up with considerations of what would have happened had the cause not occurred. As above, a cause makes a difference, and "the difference it makes must be a difference from what would have happened without it," (Lewis 1983: 557). What would have happened but did not happen is called a "counterfactual" or a situation counter to fact. "If Al Gore had been President then the U.S. would not have invaded Iraq" is a counterfactual claim. The truth conditions, or semantics, of counterfactuals are notoriously difficult to evaluate. After all, how do we know Gore would not have invaded Iraq too? We could speculate about what Gore would have done. But the question is, what could possibility make a counterfactual true?

Well, when considering counterfactuals, it is natural to think in terms of possible alternative worlds. An actual event c caused e if, had c not occurred in another possible world, then e would not have occurred in that world. And c is *not* a cause of e if, roughly, had c not occurred in another possible world, then e would still have occurred there.

Less roughly, we should note that both worlds {~c, ~e} and {~c, e} are possible worlds. What matters is which world is closer to (or more similar to, less different than) ours, the actual world. So c is a cause of e in the actual world if the possible world in which c does not occur and e does not occur is closer to the actual world than the possible world in which c does not occur but e still occurs. In yet other words, Lewis's semantics say that a causal claim's relevant counterfactual is true if e does not occur in the closest possible world in which c does not occur.[10]

Lewis asserts he is capturing the meaning of Hume's second "definition" of causation: "Where if the first object had not been, the second never existed." For an example at this stage of Lewis's theory, consider a fully functional barometer that accurately measures atmospheric pressure, which pressure causes a storm. If the barometer is working correctly then the reading we make depends counterfactually on the atmospheric pressure. That is, were the pressure to go up, then the barometer's reading would go up; and had the pressure not gone up, then the barometer's reading would not have gone up. It is also the case, of course, that if the pressure had not gone up then there would not have been a storm. So according to Lewis's theory, the storm and the barometer's going up both counterfactually depend on the atmospheric pressure, which is as it should be.[11]

---

[10] Thus, in order to evaluate a causal claim's relevant counterfactual we need a metric for deciding what counts as closeness among worlds. Lewis posits as primitive a relation of comparative overall similarity among possible worlds.

[11] Yet it appears that had the barometer not gone up then there would not have been a rise in atmospheric pressure. That's a bad result. To block this result, Lewis bars "backtracking" counterfactuals, discussed below. It also appears that the storm counterfactually depends on the barometer. Another bad result. To block this result, Lewis can appeal to his definition of causation as the ancestral of counterfactually dependent events, discussed below.

Our reconstruction of Lewis's analysis of causation is not done. As we have laid it out, counterfactual dependence among events is not necessary for causation. That is, an effect may not counterfactually depend directly on its cause. This is so in cases of so-called overdetermination. A potential but preempted cause c* (a backup cause) is an event that would have led to an effect e, except that things are such that its occurrence is blocked or preempted by the occurrence of another event c, which itself causes e.

For example, suppose that two men, Mr. White and Mr. Pink, independently of each other, are set on killing Osama bin Laden (Psillos 2002: 97). They both make similar arrangements to kill bin Laden:

> Mr. White fires his rifle; the bullet takes its course and strikes Osama in the head. Osama dies. Mr. Pink was ready to fire his rifle, and had he fired it, given his position, his shooting skills, and so on, the bullet would have also struck Osama in the head, leading to his death. But Mr. White's shot scares off Mr. Pink, who then flees the scene. Mr. Pink's shot is a potential alternative cause of Osama's death: it was preempted by Mr. White's shot, but had it not been preempted, it would have caused Mr. Smith's death.

Cases such as these demonstrate that there can be causation without counterfactual dependence because bin Laden's death is not counterfactually dependent on Mr. White's shot. It's not dependent because if Mr. White had not fired his shot then Mr. Pink would have fired his own and bin Laden would have died anyway.

The Lewisian notion of *chains* of counterfactually dependent events deals with cases of preemption like this. Causation, finally, is defined by Lewis as the ancestral of counterfactual dependence among events. There is, if we introduce plausible intermediate events, such a chain between Mr. White's shot and Osama's death, while there is not such a chain between Mr. Pink's possible shot and Osama's death.

Take an intermediate event d (e.g., that the bullet passed in between two trees) between Mr. White's shot c and bin Laden's death e. Then e is causally dependent on d and d on c. And <c,d,e> is a causal chain and it is in virtue of this chain that c caused e.

There is no such chain between Mr. Pink's shot and Osama's death; and, correctly, there is no causation there as well. A still more important point here is that even though e would still have occurred had c not occurred, there is an intermediate event d that would not still have occurred, had c not occurred. Had Mr. White not shot his gun, Mr. Pink would have shot his gun, but the bullet from Mr. Pink would have taken a different spacetime trajectory than Mr. White's bullet actually did (for example, it would have been fired a little bit later), so d would not have occurred. And had d not occurred e would not have occurred. So e counterfactually depends on d. And d on c and so there is a chain of counterfactual dependence between c and e such that c is a cause of e, even though had Mr. White not shot his gun, Mr. Pink would have killed bin Laden.

There is a problem here however. One may doubt that the effect e counterfactually depends on the intermediate event d. That is, one might think about the counterfactuals this way: if d had been absent, c would also have been absent, since d is there because c caused it to be; but then c* would have occurred and caused e. In order to block this, Lewis denies the counterfactual, "if d had been absent, c would also have been absent," as sort of "spurious reverse dependence." He calls such a thing a "backtracking counterfactual," because it would make a temporally prior event be counterfactually dependent on a temporally posterior event.[12] The past would depend on the future.

A "rich" (Hall and Paul forthcoming) and "lively area of research" (Hitchcock 2008: 323) has grown up around variations on cases such as this. The bin Laden case is a case of

[12] See Lewis (1986: 32-65).

early, asymmetrical preemption or overdetermination. In the literature, there are also cases of late asymmetrical preemption in which there are supposedly no intermediate events between cause and effect. For example, Billy and Suzy are both poised to throw their rocks at a window and cause it to shatter. Suzy throws first and her rock shatters the window. If Billy decides not to throw as soon as Suzy throws, then we have a case of early preemption like the bin Laden case. But if Billy waits to see whether Suzy's rock shatters the window, then we have a case of late preemption: There is no event between Suzy's throw and the shattering of the window on which the shattering is counterfactually dependent since Billy's throw is not preempted until the window shatters.

d. Probabilistic Causation

The program of probabilistic causation is intended to characterize causes as events that change the probability of their effects. Here is Reichenbach's early formulation:

C is a cause of E if $P(E|C) > P(E \mid \sim C)$[13]

In other words, C is a cause of E if the probability of the effect given the cause is greater than the probability of the effect absent the cause.

As we have seen, it is important that any account of causation allow us to differentiate a purported cause that is merely correlated with an effect from an actual cause of the effect. The probabilistic account can make this differentiation in its own terms. Mere correlations are discovered to be spurious when the purported cause is "screened off" by a common cause. D screens C off from E (with D causing E and C not causing E) just in case:

i. $P(E|DC) = P(E|D\sim C)$

[13] Where $P(E|C)$ is defined as $P(E \ \& \ C)/P(C)$.

ii. $P(E|{\sim}DC) = P(E|{\sim}D{\sim}C)$

This means that D screens off C just in case the probability of the effect given the screener off and the purported cause is equal to the probability of the effect given the screener off without the purported cause. The second condition is such that the probability of the effect given the purported cause absent the screener off is equal to the probability of the effect given neither the purported cause nor the screener off. And these two jointly sufficient conditions for screening off strike us as intuitive for the following reasons: if a purported cause is screened off then its purported effect would be unchanged from a situation in which the purported cause did not exist and the screener off did. And because if a purported cause is screened off then the effect is just as probable given the purported cause without the screener off as it is given neither the screener off nor the purported cause.

Consider an example. In the barometer case, the low atmospheric pressure D screens off the barometer C from the storm E i.e., makes C statistically irrelevant to E. We can understand this as follows. The probability of the storm is the same in the case in which we have the pressure with the normal barometer reading as it is in the case in which we have the pressure but the barometer fails and reads a false value. Plus, the other condition is met when the pressure D screens off the barometer C. The probability of the storm *were* there, *counter to fact*, a high atmospheric pressure with the barometer still reading low is the same as the probability of the storm were the atmospheric pressure high while the barometer reads high too. We see thereby that spurious correlations are those mere correlations in which the purported cause is screened off by the common cause. The pressure screens off the barometer as a cause of the storm, which is as it should be.

e. Process Theories

Hume challenges us to give a philosophical account of any connection between cause and effect. In order to meet Hume's challenge, process theories introduce considerable conceptual novelty. For instance, Salmon (1974) talks about so-called "processes" as primitives and dispenses with events. A process is anything that has a structure and maintains a unity even as it undergoes slight changes. Objects that persist through time, therefore, are processes. So is a propagating wave. A rolling ball. Turning gears. They are continuous: a process cannot be represented as a series of discrete events. Salmon writes, "A baseball colliding with a window would count as an event; the baseball, traveling from the bat to the window, would constitute a process," (1998: 286). The continuity of the process ultimately links up cause and effect (Salmon 1984: 156-7) in the way Hume thought impossible. Processes "are the mechanisms that propagate structure and transmit causal influence in this dynamic and changing world… they provide the ties among the various spatiotemporal parts of our universe," (Salmon 1997: 66).

We were just now listing examples of causal processes. There are also such things as pseudo causal processes. Differentiating the two is key. According to Salmon (1984), a causal process is one that transmits a "mark" and a pseudo causal process is one that does not. A mark is a local modification to the structure of a process. For example, a moving car is a causal process while the shadow the car makes is not. Let's say a process P is that which would remain uniform with respect to characteristic Q in the absence of other processes; and a process P is that for which Q would be manifest over an interval spanning both A and B (A not equal to B). Then a mark would consist of a modification from Q to Q* which has been introduced into process P by means of a single local interaction at point A and which is transmitted to B (see 1984: 148).

For example, consider the classic case of a collision between two billiard balls. Two processes, P1 with characteristic Q and P2 with characteristic R, interact at point S. It is a causal interaction and each are causal processes if P1 exhibits a modified characteristic Q* throughout an interval after S and P2 exhibits a modified characteristic R* throughout an interval after S. Salmon writes: "if two processes intersect in a manner that qualifies as a causal interaction, we may conclude that both processes are causal, for both processes have been marked (i.e., modified) in the intersection with the other and each process transmits the mark beyond the point of intersection," (1984: 174).

Objections to Salmon's mark transmission view include the fact that there are intuitively causal processes that it does not recognize) and that there are intuitively noncausal processes that it calls causal. See Dowe (2007) for a more detailed critique. Additionally, Salmon's view and views like it have a hard time accounting for mental causation and causal claims from the non-physical special sciences, such as history and economics.

Related views include Dowe's (2000) conserved quality view; Fair's (1979) physical transference theory (criticized in Dowe 2000); Aronson's (1971) "unnatural" transference theory, in which a causal process is a deviation from a "natural" process (criticized in Earman 1979); and Ehring's trope persistence theory.

f. Manipulationist and Interventionist Accounts

Once unpopular and much maligned, *manipulationist* accounts have recently been improved. New *interventionist* theories offer an exciting and promising account of causal claims. At the core of all manipulationist (including interventionist) accounts is the thought that if c is a cause of e then if one can manipulate c in the right way this should be a way of manipulating or changing e. This thought is common to the early manipulationists including Gasking (1955), Collingwood (1940) and von Wright (1971); to recent versions, including Menzies and

Price (1993); and to recent interventionist formulations, including Woodward (2003), Pearl (2000) and Spirtes, Glymour and Scheines (1993).

Woodward (2003) argues that manipulationist notions of causation abound in experimental design and the sciences generally. For example, Cook and Campbell (1979) say in their highly influential book on experimental design:

> *The paradigmatic assertion in causal relationships is that manipulation of a cause will result in the manipulation of an effect*…. Causation implies that by varying one factor I can make another vary, (36, emphasis in the original).

Consider, as well, the fact that the statistician Holland (1986) expressed his view in the slogan "no causation without manipulation." From the field of economics, Hoover (1988) says that the following "definition of cause is widely acknowledged":

> *A* causes *B* if control of *A* renders *B* controllable. A causal relation, then, is one that is invariant to interventions in *A* in the sense that if someone or something can alter the value of *A* the change in *B* follows in a predictable fashion, (173, all quotations quoted in Woodward 2003: 25).

But manipulationist accounts have not fared well in the field of philosophy, until recently. Of the early manipulationists, Woodward (2009) writes:

> … their strategy has been to take as primitive the notion of manipulation (or some related notion like agency or bringing about an outcome as a result of free action), to argue that this notion is not itself causal (or at least does not presuppose all of the features of causality the investigator is trying to analyze), and to then attempt to use this notion to construct a non-circular reductive definition of what it is for a relationship to be causal, (4).

So the idea was to stipulate as primitive some shared experience of human agency and reduce causation to some notion related to it, which itself would not be causal. However, early philosophical manipulationist accounts faced a variety of problems including circularity and

anthropocentrism. For example, we should worry about circularity when von Wright (1971) writes:

> … to think of a relation between events as causal is to think of it under the aspect of (possible) action. It is therefore true, but at the same time a little misleading to say that if *p* is a (sufficient) cause of *q*, then if I could produce *p* I could bring about *q*. For that *p* is the cause of *q*, I have endeavored to say here, means that I could bring about *q*, if I could do (so that) *p*, (74).

We may object that "producing" and "doing" are already causal notions and so should not be appealed to in a reductive account of causation. In order to dodge the circularity, von Wright may reply that he was referring to a primitive, non-causal notion of experience with human agency. But then just what causes are is unhappily married to the anthropocentric notion of our experience of agency, which is itself arguably causal anyway.

The recent interventionist program, for instance in Woodward (2003), does not depend on human agency; nor does it even attempt the reduction of causation to non-causal notions. A theory does not have to be reductive to be illuminating with respect to "how causal concepts are interconnected," Woodward (27) says.[14] Additionally, Woodward adds a counterfactual aspect that the early manipulationists did not properly countenance.

The basic idea can be put this way:

> A causes B if and only if B would change if an appropriate manipulation on A were to be carried out.

Consider the familiar example in which atmospheric pressure, A, is a common cause of B, the reading of a barometer, and S, a variable representing the occurrence or non-occurrence of a storm.

[14] Woodward does yeoman's work isolating the concepts of a total, direct and contributing cause. However, there is not the space to reconstruct the discussion.

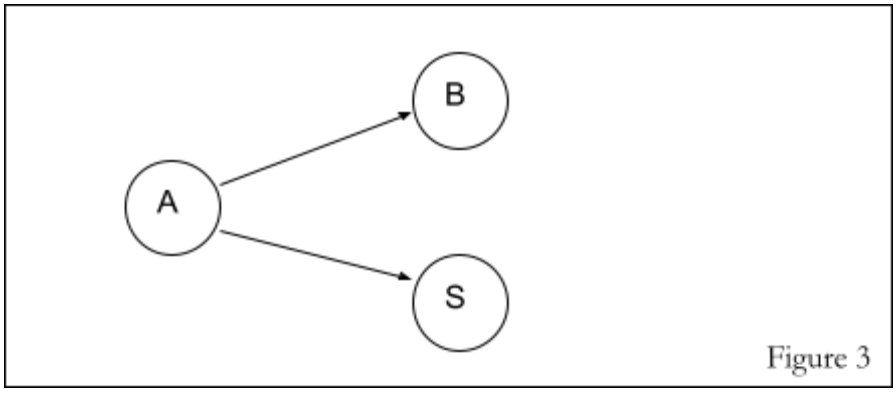

Figure 3

In Figure 3, B is represented as not being a cause of S. The interventionist account should be able to account for this in its own terms. And it does. Manipulate B directly and S does not change. Manipulate A and S indeed does change, which is the right answer.

Woodward's account can also handle familiar potential pre-emption cases. Consider again the gunmen case. "Gunman one shoots (c1) victim, causing his death, d, while gunman two does not shoot but would have shot (c2) also causing d, if c1 had not occurred," (Woodward 2009: 25).

> If we fix (via an intervention) the behavior of gunman two at its actual value (he does not shoot), then an independent intervention that alters whether gunman one shoots will alter whether victim dies, thus identifying c1 as the actual cause of e, despite the absence of counterfactual dependence (of the usual sort) between d and c1," (25-26).

Woodward's basic idea is that the claim "X causes Y" means that at least for some individuals (tokens) there is some appropriate manipulation of or intervention on the value of the variable X such that in the right conditions the value of the variable Y changes for the individual possessing it. Obviously, much depends on getting "conditions" of intervention right. We need to specify what other variables, if any, are held fixed when X is manipulated.[15]

[15] Different possibilities give different concepts of causation, including total, direct and contributing cause.

Just what counts as an intervention is "a matter of some delicacy," (Woodward 2009). We may think of an intervention on X with respect to Y as a causal process that changes X in such a way and under conditions such that any change in Y occurs *only in virtue* of Y's relationship to X.

> Heuristically, we may think of the allowable changes to X (interventions, as we have been calling them) as processes that satisfy whatever conditions must be met in an ideal experiment designed to determine whether X causes Y (Woodward 2003: 46).

Consider the barometer example again. We can randomly change the value of B by directly manipulating the dial from either high to low. "[T]he intervention 'breaks' the previously existing endogenous causal relationship between A and B" (46). But this example can now be used to illustrate more about interventions per se.

> This example illustrates the idea that interventions involve exogenous[16] changes in the variable intervened on. When an intervention occurs on B, the value of B is determined entirely by the intervention, in a way that is … independent of the value of A.

In other words, "In this sense, the interventions need to be surgical" which means the intervention does not have untoward causal consequences, i.e., "that no other causal relationships in the system are changed," (Woodward 2009: 16)

An intervention *I* on a variable *X* is always "defined with respect to a second variable *Y* (the intent being to use the notion of an intervention on *X* with respect to *Y* to characterize what it is for *X* to cause *Y*," (Woodward 2009, 17). Such an intervention must completely disrupt the causal history of *X*, as we said in the barometer example. Also *I* must not itself

[16] The interventionist theory only works on non-closed systems such that interventions can be made exogenously, i.e., from outside the system, so to speak. This has the result that the interventionist account cannot be used to model causal relationships within the entirety of the universe. Woodward replies that at the scale of the entire universe, the concept of something causing something else falls away as his theory would predict.

directly cause *Y* unless through a route that includes *X*; and *I* should not be caused by any cause that affects Y via a route that does not go through *X*. Woodward (2003) lays out other conditions in greater detail.

4. Special Issues in Causation

a. Mental Causation

When I walk into the grocery store, I like to snap my fingers just before the automatic doors open. It gives the impression that my snapping caused the doors to open, like I'm the Fonz from *Happy Days*.

Or consider this possible case: In a time of inefficient information distribution, a court astronomer could discover how to calculate when the next eclipse is going to occur – and down to the second. This would make it possible for the King to call a public assembly at that day and time and – just before the eclipse begins – announce, "In the King's name!" gesturing at the sun. This would give the impression that the King had caused the eclipse.

In each of these cases, we recognize that a purported cause is actually not a cause at all. The King's gesture doesn't cause the eclipse. Had the King done nothing, the eclipse would have occurred all the same. And my snapping my fingers doesn't cause the doors to open. It is sufficient that I walk in front of the automatic doors; they'll open whether or not I'm snapping my fingers.

Now consider a case from the philosophy of mind, specifically couched in mental terms. When a student raises her hand because she wants to make a point, we might say that her *wanting* caused the action. Yet it also seems true that said wanting – a movement of the mind, a mental event – is *realized* by a physical event in the nervous system. So we might *also* say that some state of her (complex) physical system antecedent to the action caused the

action. These two purported causes – the wanting and the physical realizer of the wanting – some philosophers say compete with respect to being the cause of the action. Then, because of orthogonal commitments to a certain kind of physicalism, it can seem to follow that *the mental event does not cause the action at all*. The purported physical cause is *sufficient* according to these physicalist commitments and "excludes" the mental event from being a cause.[17] This is the exclusion problem of mental causation; it is one of the most discussed topics in the philosophy of mind.

In the philosophy of mind literature, the exclusion problem can be represented with some relatively plausible theses that nevertheless seem to conflict with one another.

%1
%1 *Mental Causation*: Mental events cause physical events.
%1 *Non-Reduction*: Mental events neither reduce to physical events nor are identical with them.
%1 *Physical Causal Closure*: Each physical event, insofar as it has a cause, has a sufficient physical cause.
%1 *Non-Overdetermination*: Physical events are not pervasively overdetermined, i.e., they are not subject to multiple sufficient causes.

The problem is the four are inconsistent together, yet individually plausible. Anyone who holds these four theses – call them "non-reductive physicalists" – will have the exclusion problem. Most critics assume the pressure is to drop *Mental Causation*.

[17] The analogy between the automatic door, astronomy cases and the mental case breaks down at some point, of course. The first two cases offer potential causes that are not really believable while it is supposed to be a deeper quandary with respect to which one is causal in the mental causation debate, M or P. However, we might assume that the witnesses to these events believe in the causality of the finger-snap and the King's gesture. When they learned of the deceit they would "exclude" the purported causes not on their knowledge of sensor technology and astronomy but simply on an exclusion or non-overdetermination principle.

We may take a cursory look at each premise individually. We are assuming these propositions for the sake of argument. Nevertheless we can give a few considerations for or against each one.

There are considerable pressures to avoid dropping the premise of *Mental Causation* and *Non-Reduction*. The mental qua mental must be causally relevant to intentional action. We can see how important causal relevance is by observing the extent to which it is fatal to a theory if the theory entails mental epiphenomenalism.

But we should also ask a further question. Namely, why is it so important that causal relevance be explained in our account and not eliminated? *Why* is it non-negotiable? The answer lies in the practical significance of our reason explanation practices, which themselves refer to mental events as causes. Indeed, the project in the mental causation debate is to vindicate the legitimate and autonomous (thus the *Non-Reduction* thesis) explanatory status of reason explanations or explanations of intentional actions. This, admittedly, is only one part of scientific psychology and even only one part of folk psychology. But it is the recalcitrant part, the part that has resisted a naturalistic accounting.

Kim (1998) captures the non-negotiable status of the premise that some mental events cause physical events in his distinction between how and whether questions. He says the question of whether or not there is mental causation must be answered in the affirmative. Our task is to answer the question about *how* that actuality is possible.

One final way in which the first premise has been argued for has been to invoke G.E. Moore (for example in Bontly 2005). Mental causation, the fact of the causal efficacy of our mental events, is a "Moorean fact," meaning it has more in its favor than anything you could say from philosophy suggesting its non-existence. As when Moore held up his hands and said "Here is a hand and here is another," in his argument against external world

skepticism, so too can participants reply to any skeptic about mental causation, "Here is some mental causation: I am making my finger move with my mind," or some such. Mental causation is everywhere and is basic.

The third premise, again, is:

(CP) Each physical event, insofar as it has a cause, has a sufficient physical cause.

This proposition – called "the causal closure of the physical" – entails that in principle an explanation of a physical event could be accomplished without reference to any non-physical cause. It has been argued that causal closure is a key tenet of any physicalist theory. It has been argued that if any biological or psychological events are to have physical effects then they themselves must be physical events or constituted by physical events, as causal closure implies.[18]

For a long while, causal closure functioned in important arguments in the philosophy of mind without being explicitly stated. According to Papineau (2002) we need only:

> … consider J.J.C. Smart's (1958) thought that we should identify mental states with brain states, for otherwise those mental states would be "nomological danglers" which play no role in the explanation of behaviour. Or take David Lewis's (1966) and David Armstrong's (1968) argument that, since mental states are picked out by their causal roles, and since we know that physical states play these roles, mental states must be identical with those physical states.

Smart's, Lewis's and Armstrong's arguments assume something like causal closure. If something is not physical then it's a dangler without causal efficacy. This is because the only way to have effects is to be physical.[19]

---

[18] The inference however does not logically follow. "If physical then causal" is fine. But there may be other ways of being causal. However, this is not the space to challenge *Causal Closure*.

[19] Again, this does not actually follow from the way *Causal Closure* is formulated.

The fourth premise is, of course, *Non-Overdetermination*:

(NO) Physical events are not pervasively subject to multiple sufficient causes, or overdetermination.

The fourth premise is generated after contemplation of *Causal Closure*. *Causal Closure*, as it is written, does not rule out *multiple* sufficient causes. It does not exclude the effect *also* having a sufficient *mental* cause. If it also had a mental cause then we would have an event that was "overdetermined." This would be analogous to some oft-cited examples. As Kim says, "A man is shot dead by two assassins whose bullets hit him at the same time; or a building catches fire because of a short circuit in the faulty wiring and a bolt of lightning that hits the building at the same instant," (Kim 1993: 252). These are perhaps legitimate examples of overdetermination of an effect; but it is dissatisfying to imagine that a situation so statistically unlikely happens in every last case of mental causation. Every time you willed your arm to move and it moved, there would be a grand cosmic coincidence such that a neurophysiological state caused your arm to move, just at the same time as your willing caused your arm to move. Kim finds this unacceptable, and, so to disallow the possibility of overdetermination, he institutes the non-overdetermination principle.

The four plausible theses are mutually inconsistent. And that is the problem. We *could* have a consistent account if we relinquished any *one* of these, but no such relinquishing is without cost. For the sake of thoroughness, here is the logic that supports the inconsistency claim.

(1.MC) & (2.CP) & (3.NR) & entails ~(4.NO)

That is, affirming (1), (2) and (3) seems to entail denying (4), which amounts to an implausible commitment to widespread overdetermination of all mentally caused events.

(1.MC) & (2.CP) & (4.NO) entails ~(3.NR)

Affirming (1), (2) and (4) seems to entail denying (3), which amounts to mental type reduction or identity. However, as we've discussed, there are good reasons to deny type reduction or identity.

(1.MC) & (3.NR) & (4.NO) entails ~(2.CP)

Affirming (1), (3) and (4) seems to entail denying (2), which denies a key tenet of physicalism.

(3.NR) & (2.CP) & (4.NO) entails ~(1.MC)

Affirming (2), (3) and (4) seems to entail denying (1), which denies the phenomenon to be captured. Kim's version of the exclusion problem concludes that no mental event ever causes any physical event.

Notice that the problem arises whether we contrast physical properties with mental properties or with biological properties or chemical properties or geological properties. In other words, the exclusion or overdetermination problem *generalizes*. It is not really about mentality per se at all. It is about what we as physicalists are willing to accept as causally relevant.

b. Free Will and Causal Determinism

Free will and causation are mutually entangled. The "classic" problem of free will involves an apparent tension in our concepts – between our notion of ourselves as morally responsible, free agents and our notion of the universe as causally determined. The three main positions on the classic problem can be seen in how they answer this series of questions:

Is free will compatible with causal determinism?
If no, which one obtains?

Compatiblists answer the first question, "yes." Incompatiblists answer the first question, "no." The second question differentiates between two types of incompatiblists. Libertarians say free will obtains and deny causal determinism, while hard determinists say determinism obtains and deny free will. With our eye toward the role causation plays in this debate, we may notice that compatiblists and hard determinists can agree that causal determinism obtains.

A compatiblist, such as Ayer (1954), would recommend that free will not be contrasted with causation but rather with "constraint." Ayer argues the compatiblist line that every action has a cause but that does not mean every action was unfree. Instead the free actions are caused in a "special way." It seems likely we could say that the unfree actions are caused in some other special way as well. Ayer does not say more about the special way in which a free action is caused. However, he endorses the idea that when an agent acts freely she could have done otherwise, insofar as she was not constrained.

Frankfurt's (1969) widely influential paper both explicates the "special way" in which free actions are caused and provides a counterexample to the principle Ayer and other compatiblists endorse. The Principle of Alternate Possibilities (PAP) holds that "a person is morally responsible for her actions only if she could have done otherwise." By providing a counterexample to PAP, Frankfurt showed that being able to do otherwise was not as central to the problem of free will as had been thought.

The counterexample involves a man Jones who wants to kill Smith, and Black who threateningly insists that Jones kill Smith. Imagine the following possible cases. In the first case, Jones may act according to the threat: call this $Jones_1$. In $Jones_2$ the threat has no effect on him and $Jones_2$ forgets all about it and acts of his own accord, killing Smith. In $Jones_3$ the

threat strikes $Jones_3$ as worthy of consideration. $Jones_3$ acts of his own accord, but had he changed his mind then the threat would have pushed him back to killing Smith anyway. In the final situation, $Jones_4$ is not merely threatened but is made to suffer mind control at the hands of Black. Black implants a chip in $Jones_4$'s head such that if $Jones_4$ changes his mind the chip will activate and compel $Jones_4$ to kill Smith. $Jones_4$ acts on his own accord and kills Smith.

The upshot is that $Jones_3$ and $Jones_4$ killed Smith each of his own accord and in such a way that we intuitively hold them morally responsible. ($Jones_1$ is also morally responsible. Only $Jones_2$ who was compelled by the threat has diminished moral responsibility.) But while $Jones_3$ and $Jones_4$ are intuitively morally responsible, it is also the case that they could not have acted otherwise.[20] Perhaps we think $Jones_3$ could have acted otherwise by heroically resisting the threat, but it is clear that $Jones_4$ could not have acted otherwise because of the mind control chip implanted by Black. So $Jones_4$ is a counterexample to PAP. PAP says one is only responsible when one could have done otherwise, but $Jones_4$ is a case where someone is morally responsible even though they could not have done otherwise.

Frankfurt suggests that the PAP be revised to say that one is not morally responsible if one acted *only* because one could not have done otherwise. But more than that, Frankfurt's work shows, if correct, that moral responsibility does not turn on the counterfactual ability to have done otherwise. Instead responsibility turns on "the basis" of one's behavior – whether or not one "acted of one's own accord," and other *causal* concepts. If that conception could be refined then we would have an answer to Ayer's lacuna. A free action is caused in a special way, namely it is caused in a way such that the basis of one's action is one's own will or intention or values etc. In later papers Frankfurt argues roughly that a free action is one *caused*

[20] $Jones_4$ is introduced over and above $Jones_3$ because it is possible however unlikely that $Jones_3$ could resist the threat (a gun to the head). The mind control chip in $Jones_4$ cannot be resisted so $Jones_4$ offers a "cleaner" case in which one is morally responsible (for acting of one's own accord) even though one could not have done otherwise.

on a basis with which one identifies or when one's action is caused by what one "wants to want," (1974).

Bibliography

%1

%1 Ayer, A.J. 1954. *Philosophical Essays*. New York: MacMillan.

%1 Blackburn's "Hume and Thick Connexions," (1990). In his *Essays on Quasi-Realism*. Oxford: Oxford University Press, 1993.

%1 Bontly, T. 2005. "Proportionality, Causation, and Exclusion." *Philosophia* 32 (1-4): 331-348.

%1 Bromberger, S. 1966. "Why-Questions." In *Mind and Cosmos,* ed. R. Colodny, Pittsburgh: University of Pittsburgh Press.

%1 Cartwright, N. 1983. *How the Laws of Physics Lie*, Oxford: Clarendon Press.

%1 Cartwright, Nancy. 1989. *Nature' s Capacities and Their Measurement*. Oxford: Clarendon

%1 Press.

%1 Collingwood, R. 1940. *An Essay on Metaphysics*. Oxford: Clarendon Press.

%1 Cook, T. and Campbell, D. 1979. *Quasi-Experimentation: Design and Analysis Issues for Field Settings*. Boston: Houghton Mifflin Company.

%1 Craig, E. 2007. "Hume on Causality: Projectivist *and* Realist." In *The New Hume Debate: Revised Edition*. (eds.) R. Rupert and K. Richman. New York: Routledge.

%1 Davidson, D. 1967. "Causal Relations." In his *Essays on Actions and Events*, Oxford: Oxford University Press, 1980.

%1 Dijksterhuis, E.J. 1961. *The Mechanization of the World Picture*. Oxford: Oxford University Press.

%1 Eells, E. 1991. *Probabilistic Causality*, Cambridge: Cambridge University Press.

%1 Frankfurt, H. 1969. "Alternate Possibilities and Moral Responsibility," *The Journal of Philosophy*, Vol. 66, No. 23, pp. 829-839.

%1 Frankfurt, H. 1974. "Freedom of the Will and the Concept of a Person," *The Journal of Philosophy*, Vol. 68, No. 1, pp. 5-20.

%1 Gasking, D. (1955). "Causation and Recipes", *Mind*, 64, pp. 479–487.

%1 Hall, N. and L. Paul. (forthcoming). *Causation and the Counterexamples: A Travelers Guide*.

%1 Hausman, D. (1998). *Causal Asymmetries*. Cambridge: Cambridge University Press.

%1 Hempel, C. 1970. *Aspects of Scientific Explanation*. New York: Free Press.

%1 Hempel, C. and P. Oppenheim 1948. "Studies in the Logic of Explanation," *Philosophy of Science* 15 (2): pp. 135-175.

%1 Hitchcock, C. (2008). "Causation." In *The Routledge Companion to Philosophy of Science* (eds. S. Psillos and M. Curd) pp. 317-326. New York: Routledge.

%1 Holland, P. 1986. "Statistics and Causal Inference", *Journal of the American Statistical Association*, 81, pp. 945–960.

%1 Hoover, K. 1988. *The New Classical Macroeconomics*. Oxford: Basil Blackwell.

%1 Kim, J. 1998. *Mind in the Physical World*. Cambridge: MIT Press.

%1 Kitcher, P. 1989 "Explanatory Unification and the Causal Structure of the World," in *Scientific Explanation*. (eds. P. Kitcher and W. Salmon) pp. 410-505. Minneapolis: University of Minnesota Press.

%1 Leiter, B. 2002. *Nietzsche on Morality*. New York: Routledge.

%1 Lewis, D. 1973. "Causation," *Journal of Philosophy* 70 (17): pp. 556-567.

%1 Lewis, D. 1979. "Causal Explanation." In his *Philosophical Papers*, Volume II. Oxford: Oxford University Press, 1986.

%1 Mackie, J. 1973 "Causes and Conditions." In Sosa and Tooley (eds.) *Causation*. Oxford: Oxford University Press.

%1 Menzies, P. and H. Price. 1993. "Causation as a Secondary Quality." *British Journal for the Philosophy of Science,* 44: pp. 187-203.

%1 Papineau, D. 2002. *Thinking about Consciousness*. Oxford University Press.

%1 Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press.

%1 Psillos, S. 2002. *Causation and Explanation*. McGill-Queen's University Press.

%1 Salmon, W. 1971. *Statistical Explanation & Statistical Relevance*. Pittsburgh: University of Pittsburgh Press.

%1 Salmon, W. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.

%1 Salmon, W. 1987 "Probablistic Causality." In Sosa and Tooley (eds.) *Causation*. Oxford: Oxford University Press.

%1 Salmon, W. 1989. *Four Decades of Scientific Explanation*. Pittsburgh: University of Pittsburgh Press.

%1 Salmon, W. 1997. *Causation and Explanation*. Oxford: Oxford University Press.

%1 Schaffer, J. 2000. "Trumping Preemption," *Journal of Philosophy*

%1 Scheines, R. 2004. "Causation." In *New Dictionary of the History of Ideas*. Charles Scribner & Son.

%1 Smart, J.J.C. 1959. Sensations and brain processes. *Philosophical Review* 68 (April): 141-56.

%1 Spirtes, P. Glymour, C. and Scheines, R. 1983. *Causation, Prediction and Search*. New York: Springer-Verlag.

%1 Sober, E. 1985. "Two Concepts of Cause." In P. Asquith and P. Kitcher (eds.), *PSA 1984, Vol. II*, pp. 405–424. East Lansing: Philosophy of Science Association.

%1 Sober, E. 1986. "Causal Factors, Causal Inference, Causal Explanation." *Aristotelian Society Supplementary Volume* 50: 97–113.

%1 Tooley, M. 1987. *Causation: A Realist Approach*. Oxford: Oxford University Press.

%1 von Wright, G. 1971. *Explanation and Understanding*. Ithaca, New York: Cornell University Press.

%1 Woodward, J. 2003. *Making Things Happen*. Oxford: Oxford University Press.