Functionalism, Causation, and Causal Relevance

Kirk Ludwig
Department of Philosophy
University of Florida
Gainesville, FL 32611-8545
U.S.A.

kludwig@phil.ufl.edu

PSYCHE, 4(3), March 1998
http://psyche.cs.monash.edu.au/v4/psyche-4-03-ludwig.html

KEYWORDS: Functionalism, causation, causal relevance, functional explanation, consciousness, philosophy of mind.

ABSTRACT: This paper argues that functional states, and states defined in terms of them, cannot be causally relevant to the output or state transitions in terms of which those functional states are defined, or to intervening mechanisms or to anything to which their output is in turn causally relevant. Functional states therefore cannot be correctly invoked in what I call "simple causal explanations". Explanations which cite functional states are instead a species of explanation by appeal to dispositional properties. Functionalists about the mental are therefore committed either to denying that ordinary explanations of behavior by appeal to mental states are simple causal explanations or to denying that they are true. Special difficulties for functionalism arise in the case of conscious mental states.

1. Introduction

The main aim of this paper is to answer the question whether functional properties and states can be causally relevant to (a) the output in terms of which they are defined and (b) further events or states which are caused by such output. A secondary aim is to investigate the nature of explanations which invoke functional states or properties, and their relation to explanations of our behavior which invoke our beliefs, desires, and other mental states. I will argue that functional properties and states cannot be causally relevant to the output in terms of which they are defined or anything caused by that output, or to any intervening mechanism responsible for the output. The guiding idea of the argument

is that logical relations between event types precludes their standing in the causal relevance relation, which is grounded in contingent nomic relations between properties or event types. Thus, functional states, which are defined in part in terms of the output they will produce given certain input in certain circumstances, are precluded from being causally relevant to that output.<1> I will also argue, however, that this does not impugn our practices in giving explanations of the operation of systems or particular events by appeal to functional properties, though such explanations cannot be conceived of as what I will call "simple causal explanations". In light of these results, the question whether a functional analysis of mental states is adequate depends (in part) on the answer to two further questions. The first is whether ordinary explanations which invoke mental states are true. The second is whether they are simple causal explanations or what I will call "causal-functional explanations". If ordinary explanations in terms of mental states are (sometimes) true, and are simple causal explanations, then functional analyses of mental states are (in some cases) false.

The plan of the paper is as follows. In section 2, I distinguish between causation, a relation between particular events, and causal relevance, a three-place relation between event types, and circumstances, and argue for a logical independence condition on properties standing in the causal relevance relation relative to circumstances. In section 3, I apply these results to show that functionally defined states are not causally relevant to the output or state transitions in terms of which they are defined. In section 4, I extend this result to what that output in turn causes and to intervening mechanisms. In section 5, I examine the implications of this result for functional theories of mental states. In section 6, I distinguish between functional descriptions of properties and functional definitions of properties, and argue the former present no obstacle to mental states being causally relevant to behavior, but that this is so because they do not treat mental states as functional states. In section 7, I examine the nature of explanations that appeal to functional states or properties. section 8 identifies some difficulties that arise in thinking about specifically conscious mental states as functional states. section 9 is a brief conclusion.

2. Causation and Causal Relevance

Causation and causal relevance are distinct relations. Causation is an extensional relation between events. Causal relations are expressed in singular causal statements, such as

(1) My flicking the switch caused the explosion.

In (1), "my flicking the switch" and "the explosion" describe events.
A singular causal statement is of the form

(2) A caused B.

By extension, we can talk about causal relations between objects and states, or objects and states and events. Thus,

> (3) Ludwig caused the explosion,

and

> (4) Ludwig's belief that flicking the switch would cause the explosion together with his desire to do so caused it to happen.

would be parsed,

> (5) There is an event e, such that Ludwig was the agent of e and e caused the explosion,

and

> (6) There are events e1, e2, such that e1 is Ludwig's coming to have the belief that flicking the switch would cause the explosion and e2 is Ludwig's coming to have the desire to cause the explosion, and e1 and e2 together caused the explosion.

In saying causation is an extensional relation, I mean that in statements of the form (2), one can substitute in for the term that goes in for "A" or "B" any coreferring term without the truth value of the resulting sentence differing from that of the original. Thus, if the avalanche caused the flooding, and the avalanche was the most important event reported in yesterday's paper, the most important event reported in yesterday's paper caused the flooding. This shows that the causal relation is a relation between events no matter how described, i.e., it is a relation between particulars. Singular causal statements can be put in the future tense (my striking this match will cause its ignition), but typically are not put in the present tense because of the awkwardness of referring to the two events while they occur. On the other hand, we can make singular causal statements in the present progressive (my running is causing my wheezing) where we can refer to two on-going processes.

The need to distinguish between the relation of causation and causal relevance arises because in addition to an interest in what caused what, or what will cause what, we also have an interest in why one thing caused another or why one thing will cause another, and this is connected with our interest in the prediction and control of events described in specific ways. Not every description of an event describes it in a way that helps to explain why it was the cause of another event, described in a way which excites our interest. For example, the statement that the most important event reported in yesterday's newspaper caused the flooding describes one event, the flooding, in a way that makes salient what about it is of interest to us, but describes the other in a way that obscures what about it was relevant to its being a cause of the flooding. Since we often have an interest in predicting and preventing floods, we have an interest in knowing what it was

about the first event (in the circumstances) that resulted in its bringing about an event which could be described in the way salient to our interests (i.e., an event which could be described as a flooding). Since this is evidently an interest in types of events, we can, for present purposes,<2> treat the relation of causal relevance as a relation between types of events, relativized to circumstances. Relativization to circumstances is required because avalanches aren't invariably causes of flooding, but, rather, cause floods only when background conditions or circumstances are appropriate (the avalanche is of such and such a volume, in such and such a spatial relation to a river of such and such a sort, etc.). Thus, causal relevance is a three-place relation between two event types and background conditions or circumstances, conceived as a type. I will sometimes put this by talking about one event type being causally relevant to another relative to or in certain circumstances.<3>

Causal relevance interests us when we want to know why one event caused another. If we want to know why the airliner crashed, we seek, of an antecedent event, a description which meets two conditions: (i) it shows (or gives us reason to believe) the antecedent event, against the rest of the conditions present (the circumstances), to be causally relevant to the crash, and (ii) it describes an event in a way that makes it salient to our interests. (i) makes what is cited in the explanation relevant to the event to be explained. (ii) is required in addition because there will be many events which fall under types which, relative to the other background conditions present, are causally relevant to the airliner's crash (e.g., its taking off), which we are not especially interested in. These would not typically be cited in an explanation of the crash because citing them gives us information we are already in possession of, or otherwise doesn't answer to our explanatory interests, which typically involve an interest in allowing air travel, preventing crashes, and assigning blame. So we look for an explanation relevant to our interests elsewhere: inclement weather, ice on the wings, a malfunctioning engine, a bomb, a fire in the cockpit, pilot error, or the like.<4>

To this extent, the explanation of particular events is subjective. What counts as an explanation relative to one set of explanatory interests may not with respect to another. What is common to the various explanations of this type we can give of an event is that they select against the other causal conditions present a condition which satisfies (i). I will call such explanations "simple causal explanations". Such explanations are ubiquitous in everyday life, and in the applied sciences.

Causal relevance is of importance to us because it concerns the conditions for predicting and controlling events. This is a matter of discovering when certain kinds of events invariably (or perhaps mostly, or with a certain probability) follow (non-accidentally) other specific kinds of events, in particular identifiable (sorts of) conditions, or, in other words, this is a matter of discovering causal laws. One sort of event is causally relevant to another sort in certain sorts of circumstances if and only if (abbreviated "iff") there is a causal law that subsumes events described as falling under these types in the circumstances. (We can also talk about properties, state types, and object types as being causally relevant to event types, properties, state types, and object types; as in the case of the causal relation, this would be spelled out in terms of their relations to event types. I

will allow myself some license in switching back and forth to avoid awkwardness of expression.)

If we were clearer about the concept of a causal law than we are, we would be clearer about causal relevance as well. My current concern, however, is not to provide an analysis of the concept of a causal law, but to use this connection with causal relevance to bring out an important condition on two event types standing in the relation of causal relevance, relative to certain circumstances, which is often overlooked or ignored, and which is of particular importance in thinking about the role of functional states and properties in explanations. I will first discuss deterministic causal laws and then show how to generalize the account to probabilistic causal laws, of which deterministic causal laws are a special case. I will reserve "causal relevance" for "deterministic causal relevance", and speak of "probabilistic relevance" for the general case.

If a (deterministic) causal law subsumes events described as falling under types E1 and E2 relative to circumstances of type C, then it is *at least* nomically necessary that if an event of type E1 occurs in the circumstances, an event of type E2 will occur. This is a necessary condition on one event type being causally relevant to another relative to certain circumstances. However, it must also be that it is *at most* nomically necessary that if an event of type E1 occurs in the circumstances, an event of type E2 will occur. Causal laws hold in virtue of the types of events and circumstances they involve, but are contingent. No amount of conceptual analysis will tell us that heating a copper rod is followed (ceteris paribus) by its expansion, or that dropping a burning match into a keg of gunpowder will (ceteris paribus) be followed by an explosion. We establish causal laws by more or less sophisticated forms of inductive reasoning.<5> An immediate consequence of this is that predicates which appear in statements of laws must be logically independent of one another relative to the circumstances, in the sense that the satisfaction of neither in the circumstances semantically suffices for something else's satisfying the other. Otherwise, we would be able to establish a priori causally necessary or sufficient conditions for the occurrence of events of various types relative to various circumstances. Since what event types stand in the causal relevance relation relative to certain circumstances is determined by what predicates appear in true law statements involving the circumstances, event types standing in the causal relevance relation must likewise be, as I will say, logically independent of one another. This is shorthand for saying that statements that an event of the one type has occurred in the circumstances are logically independent of statements that a distinct event of the other type has occurred in the circumstances.<6> This means at least that if event types E1 and E2, relative to circumstances C, are to stand in the causal relevance relation, the statement that an event of type E1 occurred, in circumstances of type C, does not entail that a distinct event of type E2 occurred, or vice versa.

Let us consider in illustration a special class of property pairs that are not logically independent of one another, which I will call "causally committed properties" or "causally committed event types".<7> This class will be of special relevance in considering functional properties. A causally committed property is one whose instantiation requires that the event that has it cause or be caused by another event or an

event of a specific type. The property of being a kicking of a field goal (in American football), e.g., is a causally committed property, in this sense, since an event's having it entails that it caused an event of a football's clearing the crossbar between the opponent's goal posts. Similarly, the property of having jet lag is causally committed because it requires inter alia that the sufferer's condition be caused by travel across time zones by air. Causally committed properties of events (which define an event type), come in two varieties, forward looking and backward looking, which are expressed by predicates of the form:

(7) is a cause of F

and

(8) is caused by F,

respectively. Clearly, given this kind of description of an event, one could infer a priori that an event of the sort characterized by what goes in for "F" occurred. Equally clearly this would not give us any insight into what it was about the event that was causally relevant to the occurrence of an F-event. We would learn that something about the one event was causally relevant to something about the other, in the circumstances, but not what. Being told that a certain event was a kicking of a field goal does not tell us what about it was relevant to its resulting in the flight of the ball between the goal posts; we may have a pretty good idea, but not simply in virtue of being told what the event caused.

I define what it is for one event type to be *weakly logically independent* of another event type in certain circumstances as follows:

(WI) For any event types E1, E2, and circumstance type C,

E1 is *weakly logically independent* of E2 in C

iff

1) that an event of type E1 occurs in C does not entail that a distinct event of type E2 occurs;

2) that an event of type E2 occurs in C does not entail that a distinct event of type E1 occurs.

In symbols, where "E1", "E2", are restricted variables ranging over event types, "C" is a restricted variable ranging over circumstance types, "WLI(x,y,z)" stands for "x is weakly logically independent of y in z", and "ENT(x,y,z)" stands for "that an event of type x occurs in z entails that a distinct event of type y occurs":

(E1)(E2)(C)(WLI(E1,E2,C) <-> (-ENT(E1,E2,C) & -ENT(E2,E1,C))).<8>

The requirement of weak logical independence between event types in circumstances is sufficient to capture the requirement that predicates of the form of (7) and (8) don't express properties that can be causally relevant to event types characterized by what goes in for "F".

It is not strong enough, however, to rule out all pairs of event types which are excluded from the causal relevance relation. To see this, consider the property of being color-caused, defined as follows: an event has the property of being color-caused if and only if an event which caused it involved something colored. For example, the development of a photograph of someone in a red shirt is color-caused. Clearly, a shirt's being colored is not causally relevant to the development being color-caused. The requirement that event types satisfy (WI) in order to stand in the causal relevance relation correctly excludes events characterized by this pair of properties. However, consider the property of being red. Being red is logically sufficient for being colored, which is not weakly logically independent of being color-caused, but being colored is not sufficient for being red, and thus being red is weakly logically independent of being color-caused. Yet it is clear that being red is no more causally relevant to being color-caused than being colored. The reason is that the relevance of being red to being color-caused is logical, not causal, since it is logically sufficient for something logically necessary for something to be color-caused. Similar remarks apply to the property of being a color-cause, i.e., a cause of some event involving a colored object. Being red is weakly logically independent from being a color-cause, but logically sufficient for a property that is not, and, hence, cannot stand in the causal relevance relation with being a color-cause. Let us define "instantiation sufficient" as follows:

(IS) For any event types E, E+, and circumstance type C,

E+ is instantiation sufficient for E in C

iff

that an event is E+ in C entails that it is E.

Employing this definition, we can see that the examples above illustrate the following principle:

(9) For any event types E1, E2, and circumstance type C,

if E1 is not weakly logically independent of E2 in C,

then, for any event type E2+,

if E2+ is instantiation sufficient for E2 in C,
then E1 is not causally relevant in C to E2+.

(For example, instantiate "E1" to "being color caused", "E2" to "being color involving", and "E2+" to "being red involving"; (9) says that(omitting relativization to circumstance) given that being color caused is not weakly logically independent of being color involving, and being red involving is instantiation sufficient for being color involving, being color caused and being red involving do not stand in the causal relevance relation.) To satisfy (9), we can use our definitions of "weakly logically independent" and "instantiation sufficient" to define "strongly logically independent":

(SI) For any event types E1, E2, and circumstance type C,

> E1 is *strongly logically independent* of E2 in C
>
> iff
>
> 1) E1 is weakly logically independent of E2 in C;
>
> 2) for every event type E*, if E* is instantiation sufficient for E1 in C, then E* is weakly logically independent of E2 in C;
>
> 3) for every event type E*, if E* is instantiation sufficient for E2 in C, then E* is weakly logically independent of E1 in C.

In symbols, with "IS(x,y,z)" standing for "x is instantiation sufficient for y in z" and "SLI(x,y,z)" for "x is strongly logically independent of y in z", and the obvious conventions for restricted variables:

> (E1)(E2)(C)(SLI(E1,E2,C) <-> (WLI(E1,E2,C)
> & (E*)(IS(E*,E1,C) -> WLI(E*,E2,C))
> & (E*)(IS(E*,E2,C) -> WLI(E*,E1,C))))

(SI) allows us to formulate a necessary condition on two event types standing in the causal relevance relation relative to some circumstances which will meet the requirement in (9).

> CR. An event type E1 is causally relevant to an event type E2 relative to circumstances C only if E1 is strongly logically independent of E2 in C.

CR ensures that being red cannot stand in the causal relevance relation with being color-caused and or being a color-cause, as required.<9>

So far I have been assuming that the laws that underlie causal relevance relations among event types are deterministic laws. One may wish to relax this requirement in the light of the widespread acceptance that the fundamental laws of physics are indeterministic. A parallel condition, motivated in the same way as the deterministic condition, can be given readily. I define "probabilistically strongly logically independent", the condition parallel

to (SI), as above, in three stages. First, we define "probabilistically weakly logically independent":

(PWI) For any event types E1, E2, and circumstance type C,

> E1 is *probabilistically weakly logically independent* of E2 in C
>
> iff
>
> 1) that an event of type E1 occurs in C does not entail that there is an n such that n is the probability that a distinct event of type E2 will occur;
>
> 2) that an event of type E2 occurs in C does not entail that there is an n such that n is the probability that a distinct event of type E1 occurs.

Second, we define "probabilistically instantiation relevant" as follows:

(PIR) For any event types E, E+, and circumstance type C,

> E+ is probabilistically instantiation relevant to E in C
>
> iff
>
> that an event e is of type E+ in C entails that there is an n such that n is the probability that e is an event of type E in C.

Third, we employ (PWI) and (PIR) to define "probabilistically strongly logically independent".

(PSI) For any event types E1, E2, and circumstance type C,

> E1 is *probabilistically strongly logically independent* of E2 in C
>
> iff
>
> 1) E1 is probabilistically weakly logically independent of E2 in C;
>
> 2) for every event type E*, if E* is probabilistically instantiation relevant to E1 in C, then E* is probabilistically weakly logically independent of E2 in C;

3) for every event type E*, if E* is probabilistically instantiation relevant to E2 in C, then E* is probabilistically weakly logically independent of E1 in C.

In symbols, with the obvious translations,

(E1)(E2)(C)(PSL(E1,E2,C) <-> (PWL(E1,E2,C)
& (E*)(PIR(E*,E1,C) -> PWL(E*,E2,C))
& (E*)(PIR(E*,E2,C) -> PWL(E*,E1,C))))

We can formulate then a necessary condition on probabilistic causal relevance,

PCR. E1 is probabilistically causally relevant to E2 in C only if E1 is probabilistically strongly logically independent of E2 in C.

According to PCR, one event type is probabilistically causally relevant to another (in certain circumstances) only if its occurrence does not entail that there is a probability that a distinct event of the other type occurs, or entail that there is a probability that a distinct event of some type occurs whose occurrence would entail that there is a probability it is an event of the other type.

3. Functionalism and Causal Relevance

It will be immediately evident that functional states, conceived of as defined in terms of relations, causal or otherwise, to input, output, and other states of the system they are states of, cannot be causally relevant to that output or those other states.

Let us consider first functional states characterized in terms of deterministic causal relations. Any deterministic functional state can be characterized by a set of ordered triples of input, output, and future state types, . The triples represent what output O (which may be the null output) the system which has the state would produce given a certain input I (which may be the null input) and what the next state S of the system (which may be the same state) would be. A set of triples of this sort will define or partially define a functional state. A system is in that functional state provided that the set of triples correctly characterizes the counterfactual relations among input, output, and other states of the system. For the functional state to be well-defined, if the transition state is itself a functional state, it must be defined as well. This makes clear that to describe a system functionally (in the sense at issue here) is to ascribe to it a complex dispositional state, or set of interlocking dispositional states. (This makes clear also the sense in which analytic functionalism in the philosophy of mind or cognitive science, when input is characterized in terms of stimulus and output in terms of behavior in a way that does not entail anything about a system's psychological states, is a sophisticated form of logical behaviorism.)

The immediate importance of this general way of characterizing deterministic functional states is that it shows that functional states are not causally relevant to the output in terms of which they are (partially) defined. (In line with my remarks above, when I say a functional state is causally relevant to output of some sort, I mean the event (type) of the system's coming to be in that functional state is causally relevant to the event (type) which is its having that output.) The reason a functional state F is not causally relevant to output O in terms of which it is defined is that for it to be causally relevant to O it must be that there are circumstances C in which it is sufficient for O; but given the definition of F, those circumstances must include some input, I, specified in a triple which partially defines F. But if C entails that the system received input I, then C and F are logically sufficient for output of type O to be produced. F thus fails to meet a necessary condition, CR, on standing in the causal relevance relation with O. The same point, of course, applies to the relation between being in functional state F, and coming to be in any of the functional states in terms of transitions to which F is partially defined.

The same result will be reached when we turn to considering functional states defined in terms of probabilities for output given input and probabilities for transitions to new states. A probabilistically defined functional state will be characterized by a set of triples of an input, a set of probabilities for outputs which sum to one, and a set of probabilities for transitions to states which sum to one,

<I, {p(O1), ..., p(On)},{p(S1), ..., p(Sn)}>.

Deterministic functional states are a special case of probabilistically defined functional states. For a given functional state F to be probabilistically causally relevant to some output O, there must be circumstances C such that there is an n such that the probability that O obtain given F and C is n. However, given the definition of F, C must include some input I in an ordered triple that partially defines F. But if C entails that the system received input I, then C and F are logically sufficient for there to be an n such that the probability that O obtain is n. Therefore, F will fail to meet a necessary condition on probabilistic causal relevance, PCR, for being probabilistically causally relevant to O. Mutatis mutandis for any of the functional states in terms of probabilistic transitions to which F is partially defined.

4. Extension to Further Events and Intervening Mechanisms

We can extend the result to further events to which the output of functionally defined states is causally relevant to and to any internal mechanism that is causally relevant to that output by appeal to principle (T):

(T) For any events e1, e2, e3, event types E1, E3, and circumstance C,

if 1) e1 causes e2 and e2 causes e3;
2) e1 falls under E1 and e3 falls under E3;
3) E1 is causally relevant to E3 in C;

then 1) there is an event type E2 such that e2 falls under E2;
2) E1 is causally relevant to E2 in C;
3) E2 is causally relevant to E3 in C.

(T) expresses the requirement that the (probabilistic or deterministic) *causal influence* of an event of a particular type which is not the proximate cause of another event be *transmitted* through a chain of events such that for each triple of consecutive events the relation of causal relevance holds between the first and the second and the second and the third with respect to the same event type subsuming the second event. Thus, since functional states are not causally relevant to the output in terms of which they are defined, they are not causally relevant to anything to which that output in turn is causally relevant. For a similar reason functional states cannot be causally relevant to any event in a mechanism internal to the system which is causally relevant to its production of a certain output, since if it were, then there would be circumstances in which it would be causally relevant to the output, namely, the combined circumstances required for the transmission of its causal influence through the chain to the output.

5. Application to Functional Characterizations of Mental States

Functionalism as a thesis about the nature of mental states can be taken either as an empirical hypothesis ("psychofunctionalism", I will call it, following Block, 1980) or as a claim about the conceptual analysis of mental states. It is not entirely clear what the implications are for psychofunctionalism of the results of sections 3 and 4 because it is not entirely clear how the thesis, taken as an empirical hypothesis, is to be understood. I will discuss first functionalism as a conceptual thesis about the nature of mental states, and then return briefly to psychofunctionalism.

The thesis that upon analysis we can show that our concepts of mental states are concepts of functional states identifies mental properties and states with functional properties and states. It follows from this identification that any conclusions we reach about functional properties and states in general will apply to mental properties and mental states. We can conclude, therefore, that if mental states are (in this strong sense) functional states, as characterized above, they are not causally relevant to the output and state transitions in terms of which they are defined or to what the output in turn causes or to the mechanism that produces the output (even if the relations are probabilistic).

It is important to note that as a conceptual thesis, functionalism is most plausible if it identifies mental states with functional states defined in terms of the sorts of behavior we

typically invoke mental states to explain (the movement of our limbs, the issuing of sounds from our mouths, and the like), and transitions to other mental states, defined functionally in turn. This means that, for any plausible version of conceptual functionalism which identifies mental states directly with functional states, mental states are not (either probabilistically or deterministically) causally relevant to the behavior or other mental states that we typically invoke them to explain.

A functionalist theory of the mind, however, need not identify mental states directly with functional states, characterized in terms of triples relating input, output and state transitions. If we consider a complete machine table for a system, we may want to characterize what it is to have a mind as having one or another appropriate functional organization, but deny that, e.g., ordinary psychological attitudes such as belief and desire can be identified with any particular state in the machine table, since typically what behavior one produces in response to a given input is a matter not of one particular desire or belief, but of a large network of beliefs and desires (more generally degrees of belief and relative desirabilities).<10> For this reason, it would be more appropriate to identify a desire, say, to wave a hand, with a disjunction of all the states of the system in which it would be intuitively appropriate to ascribe to it a desire to wave a hand. Suppose being in state S0 is sufficient for a system to have a desire to move its hand (but not necessary), and suppose that upon input I, the appropriate output is a waving of a hand. Being in S0 is not causally relevant to the waving of a hand. But the desire to move the hand is not identical with S0. Can it, therefore, be causally relevant to the movement of the hand? No, since it is instantiation sufficient for something which is not weakly logically independent of waving a hand in the circumstances. That is, the desire is not *strongly* logically independent of waving a hand, in the circumstances. Loosening up the relation between mental states and functional states in this way is therefore no help to the functionalist in restoring the causal relevance of mental states to the output in terms of which we would expect them to be defined, and which we would invoke them to explain. Mutatis mutandis for probabilistic causal relevance. This result will generalize to any relation between functional and mental states according to which being in a given functional state is conceptually sufficient to be in a mental state.

Where does this leave psychofunctionalism? If psychofunctionalism, however conceived, allows mental states to be strongly logically independent of the behavior which we typically invoke them to explain, then it is no threat, so far as anything that has been said here goes, to the causal relevance of mental states to behavior and other mental states. On the other hand, since functional states would not be causally relevant to the behavior in terms of which they were defined, in order to maintain that mental states were causally relevant to the behavior they were ordinarily invoked to explain, some clarification would be required of the nature of the empirical hypothesis that mental states are functional states, and of its interest.

6. Functional Descriptions of Non-Functional Properties

It is important to distinguish between *functional descriptions* of states or properties and *functional definitions* of states or properties. The former identify properties or state types by descriptions of their roles in producing output and transitions to other states given certain input, but do not express the properties or types. The latter express the properties or types being defined. Functionally defined properties cannot stand in the causal relevance relation. In contrast, functional descriptions of properties do not preclude the properties described from standing in the causal relevance relation, and may require them to do so.

For example, suppose that natural kind terms, like "water" or "gold" are introduced by *functional descriptions* of the properties which they are to express. We might introduce "gold" by a functional description of the following form:

> (10) For all x, x is gold iff x has the property responsible for bright, yellow, lustrous, metallic substances having a melting point of 1063 degrees Centigrade.

In symbols, where "Gx" stands for "x is gold", "[Qx: Fx](Gx)" is the notation for restricted quantification (read: "Q x such that x is F is such that x is G"), "P" is a restricted variable ranging over properties and "R(x)" stands in for "is responsible for ...",

> (x)(Gx <-> [the P: R(P)](x has P)).

By empirical investigation, we discover, say, that the property described on the right hand side of (10) is the property of being an aggregate of an element with an atomic number of 79. This then fixes the property expressed by the predicate "is gold". Call this "G". The property expressed by the predicate on the right hand side of the above biconditional, namely,

> x has the property responsible for bright, yellow, lustrous, metallic substances having a melting point of 1063 degrees Centigrade,

is a different property. Call it "P". An object has P if and only if it has another property, namely, G. Thus, the description on the right hand side of (10) picks out a property which the predicate that employs it does not express. This is to be distinguished from defining "is gold" as expressing a functional property, as in,

> (11) Def. For all x, x is gold iff x melts when heated to 1063 degrees Centigrade.

In (11) the predicate on the right hand side of the biconditional expresses the property to be expressed by "is gold". More generally, the distinction to be drawn is between predicates introduced by means of a functional description of the property they are to express, as in,

> (12) for all x, x is F iff x has the property causally responsible for R in circumstances C,

and predicates defined as expressing a functional property, as in

> (13) Def. for all x, x is F iff x causes (or undergoes change) R in circumstances C

The arguments given in sections 4 and 5 do not show that properties or states identified by functional descriptions cannot be causally relevant to output cited in the identifying description. The property expressed by "is F" defined as in (13) cannot be causally relevant to R in C. However, the property (if any) expressed by "is F" defined as in (12) must be causally relevant to R in C.<11> I have been concerned with functional definitions of properties, not with properties picked out by functional descriptions.

The distinction between these two sorts of functional "definitions" helps to bring out more clearly why functional properties and states cannot be causally relevant to states and outputs in terms of which they are defined. Functional states are defined in terms of counterfactual conditionals without adverting to any of the laws which support the counterfactual conditionals. It is for this reason that functional states are multiply realizable, and for this reason that we can pick out causally relevant properties by means of a description of the functional roles they play in a system. It is because an event has a *specific* causally relevant property that it can cause another event's occurrence; it is because an object has *some one or other* causally relevant property that it can be in a particular functional state.

7. Functional Properties and Explanation

If the argument of sections 4-5 is correct, then functional states and properties (and states and properties defined in terms of these) cannot be causally relevant to the output or states in terms of which they are partially defined. This means that functional states and properties cannot be correctly employed in simple causal explanations of such output or transitions to such states. Since we do routinely invoke functional states and properties in causal explanations of some sort, there must be a variety of causal explanation in which functional properties and states are properly invoked. I will call these "causal-functional explanations". This general heading may, of course, cover a number of subclasses. The main purpose of this section is to sketch briefly the nature of causal-functional explanations.

Explanations are answers to questions of various sorts, most prominently why-questions, but also how-questions, what-questions, which-questions, andso on.<12> The aim of an explanation is to provide information relevant to answering a particular type of question. A question need not be asked for an explanation to be provided, for it may be evident enough what question one is concerned with. The sentence used to ask a question will often underdetermine what question was asked. For example, "Why did Eve eat the apple?" may be used to ask why Eve *ate* the apple, rather than throwing it away or adding it to her collection, or it may be used to ask why Eve ate *the apple*, rather than the pear or

passion fruit, or, perhaps, why *Eve* ate the apple, rather than Adam. The variation here has to do with the implicit contrast class the questioner has in mind, the range of alternatives to the presupposition of the question, that Eve ate the apple, relative to which he seeks information which entails or shows the presupposition of the question was to be expected. Similarly, the respect in which an answer is to explain the presupposition of the question may vary with the context, e.g., whether what is sought is an explanation of the role or purpose of something in a system or whether what is sought is explanation by subsumption under a causal law, or whether what is sought, in the case of explanation of human action, is inter alia a motive and belief in the light of which the action is revealed as reasonable for the agent. Causal-functional explanations then should be conceived of as providing information in response to a certain kind of question, in this case, a why-question. We can see what sort of why-question by looking at the answer. Of particular importance in thinking about what variety of causal explanations causal-functional explanations are will be the respect in which an explanation is sought.

As mentioned above (in section 3), to describe the functional organization of something is to say something about its dispositional properties. This gives us information about what the system will do in various circumstances. Dispositional properties are defined by counterfactual conditionals. In the characterization of functional states above, each triple corresponds to a conditional of the form: if the system were to receive input I, it would produce output O and move to state S. The characterization of an object's dispositional properties in terms of a machine table gives us information about a rich pattern of counterfactual connections between input and output for a system.

An object's dispositional properties are grounded ultimately in categorical properties, i.e., non-dispositional properties, which are the properties causally relevant to the output the object yields in response to input (for a determinate specification of a dispositional property, the input and output must themselves be characterized in terms of categorical properties rather than dispositional properties).<13> Often enough, we don't know the categorical properties of an object that ground its dispositional properties. Dispositional properties are therefore often the only properties we are in a position to cite in trying to explain why an object reacted in certain circumstances in the way it did. Explanations in terms of dispositional properties are far from vacuous, however, even in the case of very simple dispositional states. For example, for an appropriate audience, being told that salt is water soluble, and that a spoonful of salt was poured into his coffee, will explain why it tastes salty. This is explanatory, even if minimally, because it tells us that there are categorical properties that salt has that result in its having a disposition to dissolve in water, and that it is the manifestation of that disposition with respect to a spoonful of salt that is responsible for the coffee's tasting salty.

A rich source of explanations by appeal to functional states and properties of course is provided by appeals to the fact that a computer is running a particular program. Why did my computer screen go blank? Answer: I'm running a screen-saver program that blanks the screen if no key is pressed after five minutes. The explanation provides a minimal amount of information about the functional organization of my computer system, but enough in the context to answer the question given the usual explanatory interests with

which it would be asked. That is, the answer cites something that, together with what we can suppose our background information to be, entails that the screen went blank, shows therefore that it was to be expected in the circumstances, and rules out other elements of the contrast class (such as, e.g., that it is malfunctioning, or there that was a power outage). Given such an explanatory interest, it is clear that citing a fact about the functional organization of my computer is in fact sufficient to meet it. More complex functional descriptions play the same sort of explanatory role. Why is the sorting proceeding so slowly? The computer is running a bubble sort routine, which is not very efficient. The fact cited is a fact about the functional organization of the computer, but this fact about the functional organization carries information about the relative speed with which the system completes a given task in comparison with systems with different functional organizations. In none of these cases does one need know anything about the actual causal underpinnings of the system's functional organization.

If these simple examples are illustrative of the genre, we can also say something a bit more general about causal-functional explanations. There is no reason to suppose that the contrast class for causal-functional explanations is any less interest relative than it is in the case of other sorts of explanations. However, causal-functional explanations have in common the aim of citing some dispositional state or property of a system in the light of which a particular event or state (the presupposition of the often implicit question) can be seen to be what was to be expected relative to its contrast class. This is the respect in which an explanation is sought. In contrast, a request for a simple causal explanation looks for more information than does a request for a causal-functional explanation, since it seeks to identify causally relevant properties of event types.

To sum up, then, explanations serve to fill in gaps in our knowledge. Often explanations, particularly as answers to why-questions, do this by placing the fact to be explained in a pattern in the light of which it was to be expected. Some of these patterns are those formed by the network of causal laws. Simple causal explanations appeal to properties or event types connected by these laws. However, since the details of the workings of many of the objects around us remains dark to us, we often advert to explanations which cite not what ultimately underlies the complex patterns among input and output in systems we interact with but the patterns themselves. (Often enough no more than a hint is needed, since many of the patterns are familiar to us.) Citing those patterns shows that the event or state in question was to be expected, and points to an ultimate explanation in terms of the system's structure and its causally relevant properties, and tells us that, whatever they are, they are responsible for such and such systematic behavior in such and such circumstances, mediated by a perhaps complex internal causal organization.<14>

8. Conscious Mental States

The discussion so far has not distinguished between conscious mental states and others. However, it is with respect to conscious mental states that most concern arises about functional analyses of mental states. The reasons for this lie, in part, in a number of

interconnected traditional ways of thinking about conscious mental states, all of which tend to support the view that explanations of events in terms of conscious mental states are simple causal explanations.

First, there is the traditional view that the distinction between conscious mental states and others is to be conceived of as the distinction between dispositional and occurrent mental states, where occurrent mental states are manifestations of dispositional mental states. For example, if we say that someone is angry about, say, the systematic destruction of the world's rain forests, it is clear that typically what is meant is that he has a kind of dispositional state which can be attributed to him even when he is thinking about something else or is asleep. It is common to explain this sort of dispositional state as in part a disposition to experience certain emotions in appropriate conditions, e.g., upon being prompted to speak one's mind about the systematic destruction of the world's rain forests. Treating conscious mental states as manifestations of dispositional mental states, however, is incompatible with a functional analysis of them. Functional states are dispositional states, so if conscious mental states are analyzed as functional states or in terms of functional states, they cannot be treated as manifestations of dispositional states of any kind. This is clearly connected with the issue of whether conscious mental states can be thought of as being causally relevant to our behavior. If we conceive of conscious mental states as manifestations of dispositional states, then they are candidates for being subsumed by causal laws; if we think of them as dispositional states, they are not. If we cite conscious mental states in explanation of our behavior or changes in our bodies, and we conceive of them as occurrent in the above sense, then it seems clear we are treating such explanations as simple causal explanations.<15>

Closely connected with this is the view that we have non-inferential knowledge of (or warrant, perhaps defeasible, for our beliefs about) our conscious mental states. Familiar difficulties arise in trying to accommodate our having non-inferential knowledge of our conscious mental states if conscious mental states are given a functional analysis. For it is difficult to see how we could have non-inferential knowledge that we were in a dispositional state of a certain kind, since this would seem to involve knowledge of the truth of certain counterfactual conditionals, which could only be confirmed inductively. The view that conscious mental states are occurrent states is made for the view that we can have non-inferential knowledge of them. If conscious mental states are occurrent states, then it is at least open to us to have knowledge of them without appeal to induction, and so without inference from prior evidence.

Finally, the fact (if it is a fact) that we have non-inferential knowledge of our conscious mental states helps to support the case that conscious mental states are invoked in simple causal explanations. We have no non-inferential knowledge of what our conscious mental states are causally relevant to, if anything, but given that we have non-inferential knowledge of them, it can seem that we have the same kind of evidence that conscious mental states are causally relevant to certain kinds of changes in our bodies as we have that placing a hand on a hot stove is causally relevant to certain kinds of changes in our bodies. We seem, e.g., to observe regularly that blood suffuses the capillaries under the skin when, or shortly after, feeling shame or embarrassment, and it can seem that in

explaining the blush by appeal to the embarrassment we are appealing to a categorical property as causally relevant to a physical change in our bodies. It is the apparent non-inferential knowledge which is the basis for the simple induction that leads to the explanation of the blush in terms of the embarrassment. No such simple inductive inference would work for a dispositional or functionally defined property, since establishing that anything had such a property would itself require induction.

These traditional views illustrate the ease with which we fall into thinking of conscious mental states as non-dispositional, and, hence, non-functional in character, i.e., as involving categorical properties of objects. However, none of this is intended to show that conscious mental states are not functional states. The aim rather has been to identify some of the difficulties in reconciling functionalism with both these traditional views and the apparently common sense observations they are founded on. I do not say that analogs of these views could not be constructed on a functionalist account of conscious mental states, but clearly some redeployment would be required, and an explanation supplied for the contrary appearance.

9. Conclusion

To conclude, I have argued that we must distinguish between causation and causal relevance. The former is a two-place relation that holds between particular events, the latter is a three-place relation that holds between event types and circumstances. It is events falling under types that are causally relevant to one another in the circumstances that explains their standing in the causal relation. Since what event types stand in the causal relevance relation is contingent, and is established inductively, the event types that stand in the causal relevance relation must be logically independent of one another, in the sense argued for in (SI) and (PSI). These requirements are expressed in CR and PCR, the independence conditions on causal relevance and probabilistic causal relevance.

These requirements in turn show that functionally defined states cannot be causally relevant to the output in terms of which they are defined or to states in terms of transitions to which they are defined. Together with (T), the requirement on the transmission of causal influence, this shows that they cannot be causally relevant to either the mechanisms that produce the output in terms of which they are defined or to that to which that output is in turn causally relevant.

This shows that functionalism about mental states, conceived of as a conceptual thesis about mental states, has the consequence that mental states are not causally relevant to the output in terms of which they are defined, to any mechanism which produces the output, or to anything that output in turn causes, or to any of the states in terms of transitions to which they are defined. Psychofunctionalism, on the other hand, may well not lead to the view that mental states cannot be causally relevant to behavior, but it remains unclear how we are to conceive of the hypotheses and why it should be of interest if this is so.

Functional definitions are to be distinguished from functional descriptions of properties. Causally relevant properties can be identified by descriptions of their functional roles, but are not themselves functional properties. An identity theory like that of (Lewis, 1972) which identifies mental states by functional descriptions (by their theoretical roles) avoids the difficulty, but only by virtue of not being committed to mental states being functional states.

The conclusion that functional states and properties are not causally relevant to the output in terms of which they are defined, and consequently cannot be legitimately cited in simple causal explanations, does not show that they have no legitimate explanatory role to play. On the contrary, causal-functional explanations, explanations which appeal to functional states and properties, are a species of explanation by appeal to dispositional properties, and are both common and, given our limited knowledge of the causal workings of the natural world, indispensable.

The issue, then, has to be decided by attention to the nature of explanations that invoke conscious mental states. What the argument shows is that the following three propositions cannot be jointly true:

> (I) Ordinary explanations of behavior which cite mental states are simple causal explanations.
>
> (II) Some ordinary explanations of behavior which cite mental states are true.
>
> (III) Mental states are functional states.

The first two are sufficient for the falsity of the third. The burden of the defender of (III) is to show that (I) or (II) is false. Of the two, it will seem more promising to reject (I) (though not without argument, of course). Rejecting (II) throws our ordinary practices into confusion, and if the issue is conceptual functionalism, retaining (I) while rejecting (II) would require us to be deeply confused about the nature of the states which we invoke in our everyday explanations. For a defender of (I) and (II) the most obvious difficulty is to explain how properties, such as believing that one is tired or desiring to live longer, which we attribute to complex systems but not to their parts, can be conceived of as categorical properties, which is a necessary condition on treating them as candidates for standing in the causal relevance relation. It is difficult to see how this could be achieved without embracing either a version of emergentism or a view which attributed to at least some of the constituents of such systems proto-mental properties of some sort.<16>

Finally, I noted that special difficulties face an attempt to understand conscious mental states as functional states, for it requires finding a way of drawing the line between conscious and unconscious mental states that does not make the former manifestations of the latter, conceived of as dispositions, and a way of understanding how we could have non-inferential knowledge of conscious mental states. These traditional views are natural

companions to the view that explanations that cite conscious mental states are simple causal explanations.

It has not been my aim to settle the question whether functional analyses of mental states (of various sorts) can be successful. The aim rather has been to clarify the relation between this issue and questions about the sorts of explanations functional states and mental states (perhaps especially conscious mental states) can be invoked in to help make clearer where effort must be focused to settle the issue, and what the obligations are for proponents of opposing views.<17>

Notes

<1> This the line of argumentation, roughly, that logical relations between properties are incompatible with their standing in the causal relevance relation, has been anticipated by Dardis (1993), Fodor (1991), and Block (1990), who in particular argues on similar (though somewhat weaker) grounds that functionalism is likely committed to epiphenomenalism. (A related, stronger claim, that second order properties generally are inert, is made in Jackson and Pettit (1988).) This paper formulates the argument more fully and precisely, and extends the principles necessary for the argument to go through (Block's argument does not close the door for a number of reasons, among which are that it does not show why psychological state types defined in terms of disjunctions of functional state types are not causally relevant to behavior types to which they are logically connected (see section 5), and that it does not address probabilistic automata (see section 3)). An earlier version of part of the argument can be found in Ludwig (1994), as well a formulation of a screening-off condition on causal relevance. The idea that logical relations preclude *causal* relevance can be found in embryonic form earlier, and, no doubt, in some sense is already in Hume (1978). I was led originally to this line of thought (around 1986) by thinking about Davidson's defense (1980a) of action explanation as causal explanation against the charge that since reasons and actions were logically connected, they could not be causally connected (Hampshire, 1959; Kenney, 1963; Melden, 1961; Peters, 1958; Ryle, 1949; Winch, 1958). As Davidson shows, this is a confusion. Events are causally connected, their descriptions are logically related. "The cause of B caused B" could hardly be a necessary falsehood, though the descriptions are clearly logically related. But the confusion is born of failing to distinguish relations between particular events and types of events, and once this is cleared up, it can be seen that the objection embeds a correct insight. The underlying idea surely was that for a property or event type to be causally relevant (in my terminology) to another, instantiation of the one must not logically require the existence of a distinct event which instantiates the other. See section 2 for elaboration.

<2>I wish to leave open the possibility of a nominalist reduction of talk of event types. I have adopted talk of types and properties for convenience of exposition. All of the points of the paper could be recast in the formal mode without the ontological commitments, but at the cost of some circumlocution.

<3>These points are not meant to be novel or surprising, as the examples which I employ make clear. But nonetheless theyare often paid insufficient attention in discussions of causation, causal relevance, and causal explanation, so it is well to be clear about them at the outset. For further references and a recent sophisticated discussion of the distinction between causation and causal relevance, set in the framework of probabilistic causality, see Eells (1991). Davidson (1980b) contains a lucid discussion of the logical form of singular causal statements. A general introduction to the philosophical literature on the analysis of the concept of causation can be found in Sosa and Tooley (1993), along with a number of classical essays and a bibliography of important work on the subject. For an entry into the philosophical literature on mental causation in particular see Heil and Mele (1993) and Jackson (1996).

<4> The literature on explanation and scientific explanation is enormous. Entry points are provided by Kitcher and Salmon (1989) Pitt (1988) and Ruben (1993).

<5> These remarks are made from a broadly Humean and empiricist standpoint. I take more or less for granted that natural laws are contingent, and that they could have been otherwise. On this view, no necessity (beyond nomic or physical necessity) attaches, e.g., to the numbers assigned to constants of proportionality in natural laws. It is safe to say, I think, that this is the majority position among philosophers of science, but it is not uncontested. There is a different tradition, stretching back, perhaps, to Plato, which argues that causal laws are strongly necessary, and grounded in non-contingent relations among properties. While this would be the wrong place to take up this issue, the reader should be aware of the standpoint of the present paper, and that a larger debate lurks in the background. For some recent defenses of the rationalist tradition on laws, see Ellis and Lierse (1994), Fales (1990), Shoemaker (1984), Swoyer (1982). For representatives of the view that laws, though relations among universals, are contingent, see Armstrong (1978), Armstrong (1983, esp. ch. 11), Dretske (1977), Tooley (1987). Carroll (1994, appendix A) contains a critical discussion of this tradition. For a sophisticated representative of the regularity tradition, see Lewis (1973, section 3.3).

<6> The reader should not be misled by the use of "logically" here into thinking that in the sense intended two statements are logically independent of one another if and only if the sentences that express them are neither formal consequences of the other. I am concerned through out with analytic or conceptual connections.

<7> I borrow this terminology from Dardis (1993).

<8> "&", "->", "<->" and "-" are conjunction, the material conditional, the material biconditional, and negation, respectively. "(x)" is read as "for all x".

<9> Block (1990) doesn't formulate explicitly a general principle, but works mostly with examples, and argues for a position considerably weaker than mine. He does not argue that logical relations preclude nomic relations, for example. But this is because he doesn't formulate the position precisely enough to see that the examples he presents to show this don't violate the principle I give, since they involve background circumstances other than

those relevant to the production of effects in terms of which a property is defined. Where a causally committed property is causally relevant to an effect type in terms of which it is defined, it will be relative to different background conditions than those employed in the definition. (The examples Block has in mind involve agents recognizing that objects have dispositional or functional properties. There might be some doubt in these cases whether it really is the functional or dispositional properties which are causally relevant, since presumably we must infer objects have such properties from observations of their having other properties. It is hard to see why appeal to the dispositional or functional property itself would be needed in a simple causal explanation of someone's behavior. But there is no need to insist on this.)

<10> See Jeffrey (1983) for a detailed working out of the decision theoretic framework.

<11>This distinction corresponds to the distinction between David Lewis's version of the identity theory (Lewis, 1972), which employs functional descriptions of properties which need not themselves be functional properties, and analytic functionalism, which analyzes psychological terms functionally.

<12> I follow the pragmatic theory of explanation in what follows; see van Fraassen (1980, chapter 5) and Belnap and Steel (1976).

<13> The claim that all dispositional properties must be understood ultimately in terms of categorical properties, although widely endorsed, is not uncontroversial. It would be inappropriate to enter into this wider debate here (which is connected with the debate between rationalists and empiricists about laws of nature). I have argued elsewhere for the thesis in Ludwig (1996, section 2). For a recent defense of dispositional essentialism, with references to the literature pro and con, see Ellis and Lierse (1994).

<14>Jackson and Pettit (1988, 1990) make a similar point about the way explanations in terms of dispositional properties give us information relevant to our explanatory interests even though they fail to cite properties which underlie the dispositional properties. (And the point has been made earlier; Davidson, 1980a, p. 15, makes it in passing.) They refer to these explanations as "program explanations", although they are not thinking just about computer programs, but rather are taking the kinds of explanations we give by citing computer programs as an instance of a more general explanatory practice. They don't couch their point in the framework of the pragmatic theory of explanation, which provides a useful perspective from which to recognize the legitimacy of such explanations, but clearly they have the same sort of thing in mind. Their terminology differs also from mine, I think somewhat unfortunately. They draw a distinction between causal efficacy and causal relevance, but by their term "causal efficacy" they have in mind roughly the same thing I use "causal relevance" to express, whereas I would express what they express by using "causal relevance" with "causal-explanatory relevance".

<15> It is worth noting that the dispositionalist account of color properties is also in trouble if we give a functional analysis of conscious mental states. For if color properties are to be understood as dispositional properties of objects defined in terms of their

capacity to produce certain kinds of experiences in certain kinds of observers, then we must treat experiences as manifestation conditions for these dispositional states, and so not as functional states. The point extends to the entire catalog of secondary qualities and properties which have been assimilated to them.

<16>For a very quick overview of the basic positions on mind-body relations, see Ludwig (1995, section 2).

<17>I would like to thank Kevin Korb and two anonymous referees for *Psyche* for helpful comments on an earlier version of this paper.

References

Armstrong, D. (1978). *A theory of universals.* Cambridge: Cambridge University Press.

Armstrong, D. (1983). *What is a law of nature?* Cambridge: Cambridge University Press.

Belnap, N., & Steel, T. (1976). *The logic of questions and answers.* New Haven: Yale University Press.

Block, N. (1980). Troubles with functionalism. In N. Block (Ed.), *Readings in philosophical psychology*. (pp. 268-306.) Cambridge: Harvard University Press.

Block, N. (1990). Can the mind change the world? In G. Boolos (Ed.), *Meaning and method: Essays in honor of Hilary Putnam* (pp. 37-170). New York: Cambridge University Press.

Carroll, J. (1994). *Laws of nature.* New York: Cambridge University Press.

Dardis, T. (1993). Sunburn: Independence conditions on causal relevance. *Philosophy and Phenomenological Research, 53*, 577-598.

Davidson, D. (1980a). Actions, reasons, and causes. In D. Davidson (Ed.), *Essays on actions and events* (pp. 3-20). New York: Clarendon Press.

Davidson, D. (1980b). Causal relations. In D. Davidson (Ed.) *Essays on actions and events* (pp. 149-162). New York: Clarendon Press.

Dretske, F. (1977). Laws of nature. *Philosophy of Science, 44*,248-268.

Eells, E. (1991). *Probabilistic causality*. New York: Cambridge University Press.

Ellis, B., & Lierse, C. (1994). Dispositional essentialism. *Australasian Journal of Philosophy, 72*, 27-45.

Fales, E. (1990). *Causation and universals.* New York: Routledge.

Fodor, J. A. (1991). A modal argument for narrow content. *Journal of Philosophy,* 5-26.

Hampshire, S. (1959). *Thought and action.* London: Chatto and Windus.

Heil, J., & Mele, A. (Eds.) (1993). *Mental causation.* New York: Oxford University Press.

Hume, D. (1978). *A treatise of human nature.* New York: Oxford University Press.

Jackson, F. (1996). Mental causation. *Mind, 105,* 377-413.

Jackson, F., & Pettit, P. (1988). Functionalism and broad content. *Mind, 97,* 381-400.

Jackson, F., & Pettit, P. (1990). Program explanation: A general perspective, *Analysis, 50,* 107-117.

Jeffrey, R. (1983). *The logic of decision.* Chicago: University of Chicago Press.

Kenney, A. (1963). *Action, emotion and will.* London: Routledge & Kegan Paul.

Kitcher, P., & Salmon, W. C. (Eds.) (1989). *Scientific explanation, Minnesota studies in the philosophy of science, 13.* Minneapolis: University of Minnesota Press.

Lewis, D. (1972). Psychophysical and theoretical identifications. *Australasian Journal of Philosophy, 50,* 249-258.

Lewis, D. (1973). *Counterfactuals.* Cambridge: Harvard University Press.

Ludwig, K. A. (1996). Shape properties and perception. In E. Villaneuva (Ed.), *Perception* (pp. 324-350). Atascadero: Ridgeview.

Ludwig, K. A. (1995). Why the difference between quantum and classical physics is irrelevant to the mind/body problem. *Psyche, 2(16).* URL: http://psyche.cs.monash.edu.au/ v2/psyche-2-16-ludwig.html

Ludwig, K. A. (1994). Causal relevance and thought content. *Philosophical Quarterly, 44,* 334-353.

Melden, A. I. (1961). *Free action.* London: Routledge and Kegan Paul.

Peters, R. S. (1958). *The concept of motivation.* London: Routledge and Kegan Paul.

Pitt, Joseph (Ed.) (1988). *Theories of explanation.* New York: Oxford University Press.

Ruben, D.H. (1993). *Explanation, Oxford readings in philosophy.* Oxford: Oxford University Press.

Ryle, G. (1949). *The concept of mind.* New York: Barnes and Noble.

Shoemaker, S. (1984). Causality and properties. In S. Shoemaker (Ed.)*Identity, cause and mind: Philosophical essays,* (pp. 206-233). New York: Cambridge University Press.

Sosa, E., & Tooley, M. (eds.) (1993). *Causation.* Oxford: Oxford University Press.

Swoyer, C. (1982). The nature of natural laws. *Australasian Journal of Philosophy, 60,* 203-223.

Tooley, M. (1987). *Causation: A realist approach.* Oxford: Clarendon Press.

Van Fraassen, B. (1980). *The scientific image.* Oxford: Clarendon Press.

Winch, P. (1958). *The idea of a social science.* London: Routledge and Kegan Paul.