
MORAL FACTS AS CONFIGURING CAUSES

BY

TERENCE CUNEO

Abstract: The overarching aim of this essay is to argue that moral realists should be “causalists” or claim that moral facts of certain kinds are causally efficacious. To this end, I engage in two tasks. The first is to develop an account of the sense in which moral facts of certain kinds are causally efficacious. After having sketched the concept of what I call a “configuring” cause, I contend that the exercise of the moral virtues is plausibly viewed as a configuring cause. The second is to show that the causalist position I develop can withstand objections inspired by the work of Robert Audi and Jaegwon Kim.

Moral antirealists frequently object that, if moral realism were true, then moral facts would be explanatorily idle.¹ In particular, some moral antirealists complain that, if moral realism were true, then moral facts would be *causally* idle; they would not do any genuine causal explanatory work.² This complaint challenges the heart of what is perhaps the reigning orthodoxy among moral realists. The apparent orthodoxy among realists is that moral facts exist, but are in some interesting sense “natural” facts.³ It is commonly assumed, however, that natural facts are, in the paradigmatic case, causally efficacious. Accordingly, if moral facts are not causally efficacious, then they are not paradigmatic natural facts. And for those naturalists who believe that being causally efficacious is necessary and sufficient for something’s being real,⁴ the causal indolence of putative moral facts establishes that moral facts don’t exist. It is not surprising, then, that a chief concern of those who subscribe to the apparent realist orthodoxy has been to show that moral facts are causally efficacious.

An interesting feature of the recent moral realism/antirealism debate is that some prominent moral realists have agreed with the aforementioned antirealist complaint. Ronald Dworkin, Colin McGinn, Robert Audi, and Judith Jarvis Thomson, for example, have all argued that moral facts are not causally efficacious.⁵ Their claim is that the apparent orthodoxy is

false: Moral facts exist, but are not paradigmatic natural facts. Although I have considerable sympathy with this non-naturalist position, my purpose in this essay is to say a word in defense of what I shall call the “causalist” position – the view that moral facts do genuine causal explanatory work in the world. My primary aim will be to develop a view according to which moral facts of a certain kind are plausibly viewed as being causally efficacious and, thus, “natural” in one sense of this term. To this end, I engage in two tasks. The first is to develop an account of the sense in which moral facts of certain kinds are causally efficacious. After having sketched the concept of what I call a ‘configuring’ cause, I contend that the exercise of the moral virtues is plausibly viewed as a configuring cause. The second task is to argue that the causalist position I develop can withstand objections inspired by the work of Robert Audi and Jaegwon Kim. While engaging in these two tasks is not light work, there is nonetheless a sense in which the aim of this essay is fairly modest. A strategy commonly employed by philosophers who defend the causalist view is to argue that we have good reason to believe that moral realism is true because we have good reasons to believe that moral facts are causes of non-moral facts of certain kinds.⁶ While this is a perfectly acceptable strategy to employ, I am not on this occasion concerned to argue that moral realism is true. Rather, I simply wish to contend, contrary to philosophers such as Dworkin, McGinn, Thomson, and Audi, that if moral realism is true, then we have good reason to believe that moral facts of a certain kind are causally efficacious.

I. Causality and ontology: some assumptions

Let me begin by addressing two preliminary matters. First, any argument to the conclusion that moral facts are (or are not) causally efficacious presupposes some understanding of the notion of causal efficacy. Although I shall return to this topic, I want to make explicit several assumptions regarding the nature of causal efficacy that will shape the subsequent discussion.

In what follows, I shall accept a general account of causal efficacy that is fairly liberal in some senses and restrictive in others. The account is latitudinarian in the following four ways. First, I assume that entities of various kinds are causally efficacious. Thus, I assume that events, processes, facts, property instances, and so forth can figure in genuine causal explanations insofar as they are causally efficacious. Second, I assume that if an entity figures in a causal explanation by virtue of its being causally efficacious, it needn’t be the case that there is some causal “mechanism” at work, a transfer of energy or conserved quantity, or a law (strict or otherwise) that is projected. Third, I will assume that there is no sharp distinction between what is selected as the cause of an event (or fact) and the background conditions of what is selected as the cause of an event

(or fact). What we select as the cause of an event (or fact) is ultimately a pragmatic matter that is a function of our interests, purposes, and knowledge. Finally, I shall assume that there are (perhaps irreducibly) different kinds of causal efficacy. A thing can be causally efficacious insofar as it causally intervenes, contravenes, modifies, inhibits, prevents, sustains, triggers, structures, and so forth.

The account of causal efficacy with which I will be working is not liberal in the following two ways, however. First, I assume that a realist account of causal efficacy is true. Causal features are not projected on the world nor do they consist merely in the constant conjunction of two or more events. Second, while I assume that it is a necessary condition of something's being causally efficacious that it supports counterfactuals of the relevant sort (i.e., had *c* not obtained, then, *ceteris paribus*, *e* would not have obtained), I do not assume that it is a sufficient condition. So, I shall not assume that a purely counterfactual account of causal efficacy is correct.

I will not offer any arguments for the foregoing assumptions since I believe that they have been effectively defended elsewhere.⁷

The second preliminary matter I wish to address concerns the ontology of moral facts. I assume that constitutive of *moral realism of a paradigmatic sort* is the thesis that there are moral facts and that facts of this kind can be divided into at least two sub-species. I call facts of the first kind "general" moral facts and facts of the latter kind "particular" moral facts. General moral facts are ones that have the following logical form: If *x* is a token of some intention, action, etc., type *y* (e.g., murder), then *x* has some moral property *p*. The fact *that murder is wrong* and the moral norm *that wicked deeds ought to be despised* are, according to the present view, general moral facts. Particular moral facts, by contrast, are those facts that consist in a contingently existing particular's exemplifying one or another moral property at a time. The facts *that this murder is wrong* and *that Hitler's wicked deeds ought to be despised* are particular moral facts.

I draw this distinction between types of moral fact for the purpose of making clear just what sorts of moral fact are plausibly viewed as being causally efficacious. For the purposes of this paper, I will assume that, if moral facts are causally efficacious, then it is only particular moral facts that play this role. My fundamental reason for assuming this is that I doubt that unrestricted general facts of *any* sort are causes.⁸ On the assumption that this is correct, then it is also true that general moral facts are not causes.

II. *Virtues and causes*

The various arguments against the causalist view offered by moral realists such as Dworkin, McGinn, Audi, and Thomson exhibit a familiar pattern. The first stage is to argue that moral features of a certain kind are

not causally efficacious. Dworkin, for example, argues that necessary moral facts such as *that slavery is unjust* cannot be causes.⁹ Thomson and Audi contend that contingent moral facts such as *that an act-token of a certain kind is just* are likewise non-causal.¹⁰ McGinn argues, somewhat differently, that the property of goodness is not a cause.¹¹ Having shown that the selected kind of moral feature is not causally efficacious, the second stage consists in concluding that we have good reason to believe that moral features as such are causally impotent.

This strategy seems to me questionable for two reasons. In the first place, it is plausible to believe that the most defensible causalist positions are committed only to the thesis that moral facts of certain kinds are causally efficacious. The aforementioned strategy has force only if the types of moral feature that are claimed not to be causally efficacious are those that are most likely to be causes. Moreover – and this is the second point – I doubt that the objections to causalism raised by realists consider those kinds of moral feature that are the most plausible candidates for being causally efficacious. As a corrective to this strategy, I would like to put causalism in its best light by suggesting that it is the virtues in particular that are the best candidates among moral features for being causally efficacious.¹²

My argument begins with two scenarios. Here is the first:

Margaret and Sam have been colleagues for some years in the same department. Although collegial, they are not good friends. One reason for this is that Margaret is not a person who it is particularly easy to like. While highly intelligent, Margaret has the tendency to be at once a vocal and vigorous critic of the work of her colleagues and rather dismissive of their criticism of her own work. Recently, however, Sam has noticed that Margaret has seemed withdrawn. She rarely voices her views anymore in department gatherings. She regularly sits by herself in the cafeteria for lunch. This last year, moreover, has been uncharacteristically unproductive for her. While none of Margaret's other colleagues seems to take much notice of this, Sam suspects that all isn't well with Margaret. So, Sam asks Margaret to lunch. Initially, Margaret is a little stand-offish, but Sam is gently persistent and continues to check in with Margaret regularly. They develop a pattern of meeting for coffee and, after a short while, it becomes clear to Sam that, indeed, Margaret is struggling with serious personal issues. Upon learning this, Sam becomes very concerned about Margaret's well-being. Sam's concern isn't fleeting, but endures over the subsequent months – endures in spite of the fact that, on more than a few occasions, Margaret severely tries Sam's patience. The concern endures because Sam works hard at making time for Margaret, keeps her situation in mind, and puts aside some of his projects for her sake.

The question I wish to raise about this scenario is the following: Sam's colleagues are presented with signs of various sorts – signs that consist in Margaret's behaving and "carrying" herself in certain ways over a period of time. However, Sam's colleagues either fail to notice these signs, or if they do notice them, their noticing them does not motivate them to help Margaret. With Sam it is different. Sam interprets these signs as calling for a certain range of responses and is motivated in an appropriate fashion by virtue of being aware of these signs. That is to say, by virtue of Sam's being aware of these signs, a variety of motivational states with appropriate content and strength that endure in different contexts over an appropriate stretch of time form in the right fashion. Why?

In a moment, I'll return to this scenario. Let's now have the second scenario before us:

Both a scholar of obscure languages and a computer aficionado, Sam has several different types of computers in his office. Today he has brought Lyle to his office. Lyle is the son of American expatriates who live in an obscure village in the South Pacific. Since Lyle has been raised in a very isolated environment, he knows almost nothing about computers and how they work. As Lyle presses on certain keys on the keyboard of one of Sam's computers, letters of certain types appear on that computer's monitor screen. When he presses keys of the same type on a keyboard joined to another computer, letters from a different language appear on that computer's monitor. Lyle finds this puzzling.

Suppose we grant that the fact that Lyle pushes down on certain keys of a given computer is (to use Fred Dretske's term) a "triggering" cause of the fact that letters of certain types appear on that computer's monitor.¹³ Had Lyle not pressed on these keys, then (all other things being equal) these letters would not have appeared on that computer's monitor. Given Lyle's puzzlement, however, there is a natural question to raise: When Lyle presses down on the keys of a keyboard joined to a particular computer, letters of certain types are displayed on that computer's monitor. And when he presses down on keys of the same type on a keyboard joined to a different computer, letters of different types are displayed on that computer's monitor. Why?

I am going to contend that there is a kind of causal explanation that answers both questions I have raised.

Begin with the second scenario. Our question in this case is what explains the fact that when Lyle presses down on the same types of keys on keyboards joined to different computers, letters of different types appear on those computers' monitors. The natural answer to offer is that it is certain features of the hardware conditions (the actual electrical connections in the computer) and programming (the software) of each

computer working in a particular way that is causally responsible for this. The fact that the hardware and software components of each computer are working in a certain way “configures” the informational content introduced by Lyle’s pressing certain keys on a given keyboard in such a way so as to produce letters of a certain type on that computer’s monitor.¹⁴ We could put the matter thus: The fact that certain features of the hardware and software components of a given computer are manifested in a certain way is a configuring cause of the fact that letters of certain types appear on that computer’s monitor.

So, take an information-processing system, *S*, an event or fact *A* that is an input into *S*, and an event or fact *B* that is an output of *S*. Assume that *A* carries information of a certain kind and is a triggering cause of *B*. Assume also that there is a type of causal process *C* occasioned by *A* that mediates *A* and *B* and that transmits the information carried by *A*. Assume, furthermore, that there is no sense in which *C* had to obtain. Given *A*, any number of subsequent types of causal process could have taken place. A configuring cause is a (concrete) fact or event *F* that configures (or structures) the information carried in *A* in such a way that *C* obtains and, thus, *A* causes *B*. Or as I shall say, given *A*, *F* is a configuring cause of *B*. In Lyle’s case, the fact that he pressed on certain keys of a keyboard joined to a particular computer and the fact that letters of certain types appeared on that computer’s monitor is mediated by a causal process. Had certain elements of the hardware and software components not operated in a particular way, then (all other things being equal) the fact that Lyle pressed certain keys on that keyboard would not have caused letters of certain types to appear on that computer’s monitor.¹⁵

Turn now to the scenario involving Sam and Margaret. In this instance, we asked what accounts for the fact that Sam is motivated in a variety of appropriate ways – or as I will simply say, motivated appropriately. An appealing suggestion is that what accounts for this fact is that Sam manifests various features of the virtue of compassion. If we think of a person’s virtue as being comprised of “aspects” or capacities, we can say that Sam’s behavior is the product of the manifestation of the various aspects or capacities that comprise Sam’s compassion.¹⁶ In noticing the presence (and absence) of certain features of Margaret’s countenance and behavior, for example, Sam exercises what we might call the “perceptual” aspect of a virtue. In interpreting these features as morally relevant, and fixing upon an appropriate course of response, Sam exercises what we can call the “interpretive” aspect of a virtue. And in being moved to act appropriately in response to these features, what we can call the “motivational” aspect of Sam’s virtue is actualized. Sam’s behavior, then, is plausibly thought of as the product of the coordinated manifestation of various aspects of the particular way in which he instantiates the virtue of compassion.

According to the view I am elaborating, a virtue is a constellation of aspects or capacities. More exactly, it is a *stable* and *enduring* collection of aspects or capacities that ordinarily work in concert. What I should add is that it is plausible to believe that a virtue is not a higher-order property “over and above” this unified and stable constellation of capacities. There’s not the unified and stable constellation of the perceptual, interpretive, and motivational capacities, and then Sam’s compassion over and above that. Nor is there the manifestation of these aspects, and the higher-order property of *Sam’s acting compassionately* over and above the manifestation of these capacities. We might call this a “token identity” or “constitution” view of a person’s virtue and its manifestations – though I won’t put much weight on which label we use. What is important to see is that the view isn’t in any interesting sense reductionistic. By claiming that a person’s virtue and its manifestations are “nothing over and above” a stable and enduring collection of capacities and their manifestations, it is not as if we have reduced a person’s virtue or its manifestations to a cluster of mere naturalistic descriptive properties. A person’s virtue consists in the capacities to notice what is morally *salient*, to interpret these signs *correctly*, and to be motivated *appropriately*. A person’s virtue, then, is a richly evaluative thing that consists in being reliably disposed to respond to reasons of various kinds appropriately.

The thesis I propose is that the mutual manifestation of the interpretive and motivational aspects of Sam’s compassion is a configuring cause of the fact that he is motivated appropriately over a certain period of time.¹⁷

Suppose we assume the fact that Sam notices certain features of Margaret’s behavior and countenance is a triggering cause of the fact that Sam is motivated appropriately. Suppose, moreover, this triggering cause introduces informational or “intuitional” content about Margaret (e.g., that Margaret appears in such-and-such way, etc.) into Sam’s cognitive system. Suppose, also, there is a causal process occasioned by the fact that Sam notices certain features of Margaret’s behavior and countenance that mediates between this fact and the fact that Sam is motivated appropriately. And suppose that this process did not have to obtain; any number of different causal processes could have occurred. Finally, suppose that this causal process is comprised of mental states that themselves bear intuitional content and that these states stand in causal relations of various types to one another. The exercise of the interpretive and motivational aspects of Sam’s virtue, I suggest, “configures” the intuitional content of Sam’s awareness of Margaret’s behavior and countenance in such a way that Sam’s awareness of these signs yields the formation (and maintenance) of a variety of appropriate motivational states. In so causally configuring this information, the exercise of these aspects of Sam’s virtue doesn’t merely transmit the informational content of Sam’s awareness; it *interprets* it as being indicative of certain features of Sam’s environment.

Had these aspects of Sam's virtue not been manifested in certain ways, then (all other things being equal) the variety of motivational states that consist in Sam's being appropriately concerned for Margaret's well-being would not have been formed (and maintained).

Let me now lay bare two assumptions concerning this thesis – one that I will not defend at this point, the other which I will defend in a moment. The assumption that I will not defend at this point is the claim that there are such things as configuring causes. At this stage in the discussion, I simply recommend this assumption as a plausible claim about which a realist does not have reason to be particularly suspicious. The assumption that I will defend shortly is that there is a configuring cause of the fact that Sam is motivated appropriately. What I shall now defend is the conditional claim that, if there is a configuring cause of the fact that Sam is motivated appropriately, it is the exercise of Sam's virtue.

III. Developing the view

What we are trying to explain is a fairly unusual phenomenon. After all, among ordinary agents it is a regular occurrence that, given certain kinds of experience, they feel great concern only to have that concern evaporate shortly thereafter. It is also a regular occurrence among such agents that they fail to display appropriately compassionate behavior across different contexts. Sam's concern for Margaret, however, exhibits none of these defects. Given Sam's awareness of Margaret's countenance and behavior, he is motivated appropriately over a significant stretch of time in different contexts – even in the teeth of considerations that would threaten to extinguish this motivation. It is this fairly unusual state of affairs for which we would like to find a configuring causal explanation. Here I should like to adduce a pair of considerations that speak in favor of the thesis I propose.

In the first place, it appears to be part of the very job description of the exercise of a virtue that it explains constancies in behavior across varying circumstances, resistance to distorting influences, and the like.¹⁸ Since this is precisely what we are trying to explain, the fact that Sam has exercised the virtue of compassion is an excellent candidate for explaining the fact that he is motivated appropriately. Of course this alone doesn't give us reason to believe that, if there were a configuring cause of Sam's motivational state, then it is the manifestation of Sam's virtue. But, secondly, there are good reasons for affirming this latter claim as well. If what I've said is correct, paradigmatic configuring cause explanations are ones that appeal to the way in which an entity's "programming" configures informational content of certain kinds. To be a virtuous agent, however, arguably just is to be "programmed" in a certain fashion. The virtuous

person is programmed in such a way that, given certain kinds of experiential inputs (and when all goes well), the manifestation of certain features of that programming configures the intuitional content of those experiential inputs so as to yield mental states and activities of certain kinds. Given the fact that Sam is programmed to interpret signs of certain kinds aright, when all goes well, his noticing those signs yields motivational states and activities of the appropriate types.

If this is right, being a configuring cause also appears to be part of the job description of the exercise of a virtue. However, the causalist presumably needs to say more than this in defense of her view. What the causalist presumably wants to claim is not merely that it is part of the job description of the exercise of Sam's compassion that it is a configuring cause. Rather, what the causalist also wishes to claim is that it is in virtue of its *normative* character that the exercise of Sam's compassion is a configuring cause. That is, the causalist presumably wants to maintain that the fact that the manifestation of Sam's virtue is an *appropriate* response to reasons is fundamental to its being a configuring cause.¹⁹ What can be said in favor of this further claim?

I propose to address this question by considering an objection to the causalist position articulated by Robert Audi.²⁰ Audi's objection hinges on the claim that it is not moral facts, but the naturalistic/non-moral features that determine them that do all the causal work in putative moral explanations. The causalist reply to this objection should throw light on the sense in which it is the normative character of the exercise of Sam's virtue that functions as a configuring cause.

IV. Audi's objection

Audi's objection proceeds on a pair of assumptions. First of all, says Audi, moral facts *ontologically depend* on "natural" facts. For any particular moral fact, there will be a range of natural facts on which that moral fact supervenes and with which that moral fact is not identical. Second, our awareness of moral facts *epistemically depends* on these natural facts. That is to say, being aware of a particular's having a moral property typically depends on being aware of its having a certain range of base natural properties that determine its having that moral property. This latter type of dependency, writes Audi, has ramifications for moral explanations:

Whenever we explanatorily invoke a moral property, it will be in part on the basis of, or at least in the light of, some belief or presupposition to the effect that one or more natural properties is playing an explanatory role. We are thus in a position to rely – often unself-consciously, for sure – on those other properties to do the explanatory work, and it is arguable that they, and not any moral property, are in fact what does it. . . . Our understanding

of how the ascription of moral properties can explain (at least so far as causal explanation goes) seems wholly *derivative* from our understanding of how the relevant base properties can do so.²¹

The idea, then, is that for any putative causal explanation of a phenomenon E that appeals to the causal efficacy of some moral fact M, we can replace, without loss of explanatory power, the reference to M with the base natural facts N on which M supervenes. So, the idea is not that we can offer different *complementary* causal explanations of E in terms of M and N. Rather, it is that any putative causal explanation of E offered in terms of M can be *eliminated* in favor of an explanation in terms of N. But if this is right, then we have no need to claim that moral facts do any causal explanatory work. All the causal explanatory work is done by the natural facts on which these moral features supervene.

To see the thrust of Audi's objection more clearly, let's consider an example he furnishes. According to the example, the citizens of a particular country revolt because of governmental injustice. Audi writes of this putative explanation:

One cannot know (and normally would not even believe) that there is such injustice except through some kind of awareness of, say, government seizure of land, arbitrary curfews, and police brutality, where these are construed behaviorally in terms of, for example, soldiers' occupying farmland, clearing streets at night, and clubbing non-protesters. But these are just the sorts of non-moral factors, that in their own right, we suppose . . . can perfectly well explain a revolt. They also seem to have causal power in a quite intuitive sense.²²

Audi's view, then, is that our awareness of governmental injustice in the case described depends on both our awareness of the behavior of government employees and their intentions. We can, for instance, understand the police brutality in terms of their periodically clubbing and intending to severely harm some of the citizens. We can understand the administrative deceit in terms of the administration's making various statements intended to lead the citizens into believing falsehoods. It is these behavioral and non-moral "social-psychological" features of the world that are supposed to do the causal explanatory work in Audi's imagined revolt.²³ The injustice of the government's actions is simply epiphenomenal with respect to causally explaining the citizens' revolt.

What Audi says seems to me to present a powerful objection to causalist views of certain kinds, provided that the putatively social-psychological facts to which he adverts are themselves genuinely naturalistic/non-moral ones. While I think there is excellent reason to doubt this, I shall not pursue the point here. Instead I want to emphasize that there is an important difference between the scenario Audi sketches and the one that I have employed. According to Audi's scenario, what we are attempting to explain causally is why a non-moral state of affairs obtains, viz., *that the*

citizens revolt. By contrast, in the case with which I have been working, what we are trying to explain causally is why a moral state of affairs of a certain kind obtains, viz., *that Sam is motivated appropriately*. This difference between the two cases is important because even if Audi were right to claim that we can handily eliminate reference to moral facts when causally explaining why non-moral facts of certain kinds obtain, it wouldn't follow that we can do the same when causally explaining why moral facts of certain kinds obtain. So, even if Audi's objection were to show that moral facts do not causally explain non-moral ones, it wouldn't follow that moral facts *as such* are causally inefficacious. The question, then, that needs to be raised about Sam's case is this: Can we fashion an adequate configuring cause explanation of the moral fact that Sam is motivated appropriately that makes no essential reference to moral facts?

Suppose we pursue this issue by sketching how a rival, non-normative configuring cause explanation of Sam's motivational state might go. One such explanation, broadly Humean in character, runs something like the following. Assume, first of all, that Sam is compassionate and that his being compassionate is constituted (in part) by the fact that he is deeply and firmly concerned about the welfare of his colleagues. Assume, furthermore, that the fact that Sam is aware of signs of certain kinds is a triggering cause that occasions a causal process that itself eventuates in the fact that Sam is motivated appropriately. However, maintain that we can explain why this awareness generates an appropriate motivation by appeal to a constellation of non-moral facts that constitute the exercise of Sam's virtue. Roughly, the suggestion is that the intuitional content of Sam's awareness is configured by the manifestation of Sam's reliable disposition to form true moral and non-moral beliefs of certain types and his deep concern for his colleagues. These beliefs and concern hook up to form the appropriate motivational state. Most importantly, the beliefs and concern that determine this motivational state are not in any sense moral entities. Rather, they are simply non-normative features that both appear to have causal power in their own right and to be capable of doing all the requisite causal explanatory work to explain why Sam is motivated in the way he is.

As I've indicated, this rival configuring cause explanation is stated in fairly rough terms. I want now to argue, however, that it has to be stated thus, for once we start filling in the details, it becomes very difficult to fashion an adequate configuring cause explanation that is genuinely non-normative in character. Let me highlight three points.

First, the advocate of this rival style of explanation wants to explain the fact that Sam is motivated appropriately by appealing simply to (i) the manifestation of Sam's disposition to form true moral and non-moral beliefs of certain types given certain kinds of experiential input and (ii) the fact that these beliefs hook up with the relevant concern. But of

this putative explanation one wants to ask: What *types* of true beliefs are supposed to be adequate to the explanatory task? A person can, after all, form either too few or too many true beliefs in a given situation. Form too few true beliefs and his grip on the situation will be inadequate; he will have too little accurate information concerning the case at hand. Form too many and he will have an excess of information that obfuscates the important issues; as a result, he may well be indecisive or motivationally paralyzed.²⁴ What we're trying to explain (in part) is why Sam reliably forms and maintains a sufficiently wide range of *relevant* true beliefs and, in addition, holds their content with an appropriate degree of confidence. Merely appealing to the manifestation of a disposition to form true beliefs of certain kinds, however, won't explain that.

Second, if Sam's motivational state is a manifestation of a virtue, it matters *how* Sam forms these beliefs and how they hook up with the relevant concern. As recent discussions concerning so-called deviant causal chains have made evident, simply manifesting reliable dispositions of the types mentioned above isn't sufficient to establish that the beliefs and motivational states that are the product of such dispositions are virtuous or, for that matter, rational in any robust sense. The problem is that dispositions to form true beliefs and to be motivated in certain ways can be accidentally reliable insofar as they are the product of serious and systematic cognitive malfunction, bad habits of thinking, the meddling of cognitive scientists, and so on.²⁵ Accordingly, the mere fact that Sam's motivational states are the product of reliable tendencies to form true moral beliefs and be motivated in certain ways does not explain why those states are virtuous or rational in any full-bodied sense (we can imagine cases in which Sam's moral beliefs reliably track the truth, but are the product of unreflectively appropriating the beliefs of others, for example.) Appealing to normative configuring causes, however, will. For suppose we assume – as I have thus far – that different mental states are linked to each other by causal processes and that there is no sense in which these processes must obtain. Suppose, further, that we adopt the realist assumption that the states that result from these processes exhibit normative merits and demerits such as being rational, unjustified, conscientious, careless, and so forth. Suppose, finally, that in order for these resultant mental states to exhibit these normative features, they must have a certain kind of causal ancestry or, as I have been saying, be causally “configured” in certain ways. If all this is right, then normative facts would appear to have an important theoretical role to play insofar as they can help dissolve the problems of accidental reliability and deviant causal chains that have dogged reliabilist theories of justification and functionalist theories of the mind. All too briefly, inasmuch as normative facts of certain kinds configure causal processes that mediate different mental states – configure them in such a way that the appropriate mental states are formed or

maintained – they are (in part) what prevent aberrations of the types mentioned above from occurring and ensure that mental states are formed or maintained in a virtuous or rational fashion.²⁶ Accordingly, if the causal history of the formation of a mental state determines (in part) why that mental state exhibits certain normative merits or demerits, then normative configuring causes can help explain why.

Third, and finally, it is doubtful that Sam's case can be adequately characterized in the Humean style described above because in Sam's case there is an ineliminable element of interpretation that is also present. This element of interpretation minimally involves fitting the information present into a broader situation, balancing it with other important factors, "walling-off" certain considerations, and seeing patterns of particular kinds with a certain degree of vividness. Fundamental to the appropriateness of Sam's motivation, then, is the fact that he forms appropriate beliefs by way of interpreting his situation aright. But it is difficult to see how, in purely non-normative terms, we can fashion a configuring cause explanation of the fact that Sam interprets his situation aright, and thereby arrives at appropriate beliefs. This is not simply because the natural facts that determine the appropriateness of his interpretation form such a wildly heterogeneous group that we need to lump them under a normative concept. It is also because that among the facts that appear to configure the intuitional content of Sam's awareness is the manifestation of Sam's disposition to conduct himself *conscientiously*. However, it is difficult to think of a way in which we can give an account of Sam's conscientious conduct that does not appeal to the fact that Sam has conducted himself as he ought. Nor do I see how, having appealed to that fact, we can somehow isolate the naturalistic/non-normative features of Sam's conscientious conduct and claim that these features alone configure the intuitional content of Sam's awareness. It is not the mere fact that Sam has paid close attention to Margaret's behavior that accounts for the fact that his awareness yields an appropriate motivation. Rather, it is the fact that he has paid close attention to what (and when) he ought. And I think that once we see the intimate manner in which the normative and the natural are intertwined in "thick" moral facts such as *that Sam has conducted himself conscientiously*, we can't simply take it for granted that we can neatly sever the normative aspects of these facts from their natural features and claim that the former are mere causal idlers while the latter shoulder the real causal burden.

To this last point there is perhaps a natural response. The natural response is to dig down deeper yet and attempt to isolate the natural facts on which Sam's conscientiousness and the moral reasons to which he responds supervene. The thought is we can claim that it is *these* natural facts (in tandem with other such facts) that causally explain the fact that Sam is motivated appropriately.

While this response may be natural, it seems to me unpromising, the main problem being that it is not obvious what kinds of naturalistic/non-moral fact Sam's conscientiousness and the reasons to which he responds ontologically depend on. Exactly what natural facts make it the case that Margaret's uttering a certain string of phonemes on a given occasion gives Sam moral reasons of various kinds to pay close attention to what she says? Exactly what natural facts make it the case that Sam has exercised moral conscientiousness on a given occasion? And exactly what natural facts make it the case that Sam holds his beliefs concerning Margaret's situation with an appropriate level of confidence and that his concern for her well-being has endured for an appropriate length of time? These are extremely difficult questions to answer. And the difficulty of these questions leads one to suspect that though instances of obligations, entitlements, virtues, and the like are perhaps ontologically and epistemically dependent on various naturalistic/non-moral properties, the ontological and epistemic dependence is in many respects not *epistemically transparent*. So, it is not as if the advocate of the present response can simply pick out the natural facts that determine the fact that Sam has behaved in a conscientiousness fashion and then claim that it is these natural facts alone that causally explain the occurrence of Sam's motivational state (or even the instantiation of the properties upon which this state supervenes). It is just not obvious what these natural facts are.

To sum up: What made Audi's objection appear to be a serious threat to the causalist position is its claim that the reference to moral facts in causal explanations could be replaced without loss of explanatory power with reference to the naturalistic/non-moral facts upon which they supervene. As I've pointed out, this objection assumes that, when furnishing these rival explanations, we can isolate these subvening non-moral features well enough so as to construct causal explanations in terms of them. I've argued that this assumption is dubious. Even if we grant that moral facts are determined by natural ones, both the "shapelessness" and "thickness" of moral facts of certain types guarantees that we cannot construct the requisite type of configuring cause explanations only in terms of the determining natural features. The shapelessness of moral facts presents problems for Audi's objection because we are not able to isolate those features that determine moral ones without employing moral concepts and, thus, are unable to construct configuring cause explanations entirely in terms of these determining features.²⁷ The thickness of moral facts presents an analogous difficulty insofar as these facts present an interesting case in which the normative and the non-normative are so entangled that there is no principled reason for claiming it is the non-normative features of thick moral facts, and not their moral ones, that do all the causal explanatory work.

Let me add to this summary a final point. I have said that we cannot offer a configuring cause explanation of why Sam is motivated appropriately

by appealing simply to the naturalistic/non-moral facts upon which the exercise of Sam's virtue supervenes. However, I want also to suggest that, if what I've argued is correct, then we also have good reason to believe that we cannot offer a configuring cause explanation of the naturalistic/*non*-moral facts upon which the exercise of Sam's supervenes that appeals only to other naturalistic/*non*-moral facts. For take a pair of naturalistic/*non*-moral facts N and N* and a pair of moral facts M and M* and assume that M supervenes on (i.e. ontologically depends on) N and M* on N*. Assume, further, that N (or some subset of features that constitute N) causes N*. Finally, assume that M* is identical with the fact that Sam is motivated appropriately. If M* supervenes on N*, however, then N* must be constituted by the types of feature sufficient for its determining M*. Presumably, among these features is N*'s having the right sort of causal ancestry; if N* were the product of a deviant causal chain, for example, then it would not determine the fact that Sam is motivated appropriately. But if what I've argued is right, given the myriad of possible causal ancestries that N* might have, we're simply not in a position to isolate adequately the facts that causally configure N* without recourse to normative concepts of various kinds. Accordingly, we are in no better position to furnish a configuring cause explanation of N* that appeals only to non-normative features than we are of providing such an explanation of M*. It follows that although the causalist response to Audi's objection assumes that moral facts can be the explananda in causal explanations, it needn't; one can also employ the response to address a position that assumes otherwise.

V. *Engaging a suspicion*

Audi's objection is driven by the suspicion that putative moral explanations are not to be taken at face value. And while I've argued that the way in which Audi develops this suspicion fails to provide reason to believe that moral facts are causally inefficacious, I conjecture that the foregoing argument will not have shaken the conviction that it is only the underlying naturalistic/*non*-moral facts that are doing the genuine causal work in putative moral explanations. The reason for my conjecture is this: Audi's objection hinges on the claim – call it the “epistemic claim” – that we can consistently isolate the underlying naturalistic features doing the genuine causal work in putative moral explanations. But, it may be pointed out, causal explanations need not appeal only to causally efficacious features. So, even if the epistemic claim is false and we are forced to appeal to moral concepts in order to causally explain the existence of moral facts of certain kinds, it doesn't follow that the extension of these concepts does any genuine causal work. Accordingly, even if the foregoing reply to Audi

is on the mark, it doesn't follow that causalism is true. Moreover, it might be continued, developing Audi's suspicion needn't involve appealing to the epistemic claim. All that is needed to vindicate the suspicion is *good reason* to believe that it is simply the naturalistic features (whatever they may be) that underlie moral facts that provide the real causal muscle in putative moral explanations.

Let me address these two concerns in order.

The appropriate response to the first concern, I think, is to deny that the considerations adduced in favor of the causalist view and against Audi's objection are supposed to entail that the exercise of Sam's virtue is a configuring cause. Rather, the claim is the following: Given moral realism, the thesis that the exercise of Sam's virtue is a configuring cause *better* explains certain facts than a view that renders the exercise of his virtue epiphenomenal. The explanation is better in part because a moral realist of a paradigmatic sort is committed to the existence of the virtues and it appears part of the very job description of the exercise of a virtue that it is a configuring cause. As I've emphasized, fundamental to our understanding of the moral virtues is the claim that their exercise sustains desires, beliefs, and intentions of certain appropriate kinds, inhibits acting on certain temptations, and influences the formation and maintenance of propositional attitudes and motivational states of appropriate types – where “sustaining,” “inhibiting,” and “influencing” are all causal terms that can be understood in light of the concept of a configuring cause. It is worth adding, moreover, that by thus appealing to the apparent functional profile of the virtues, the causalist is in good company. After all, our reason for believing that the exercise of the virtues is causally efficacious is really no different from our reason for believing that the propositional attitudes themselves are causally efficacious. It is precisely because it appears to be constitutive of the explanatory profile of propositional attitudes that they are causes that we take them to be causes. For example, it appears to be constitutive of the explanatory profile of desires and beliefs that such attitudes influence behavior by sustaining and inhibiting the formation of other attitudes, influencing the formation of motives, and so forth. Accordingly, to deny that the exercise of the virtues or the propositional attitudes plays these causal roles is perforce to deny that our ordinary understanding of their job description is accurate.

Fundamental to the causalist view, then, is the conviction that, in the absence of sufficient reason to believe otherwise, the best option for the moral realist is to assume that the apparent functional profile of the virtues gives us genuine insight into their nature. And since the apparent functional profile of the virtues tells us that their exercise does genuine causal explanatory work, we should (all other things being equal) conclude that the exercise of the virtues is causally efficacious. To transpose what Tyler Burge says about mental causation, our best guide to moral causation lies in understanding our best

means of explaining why agents of certain kinds are motivated appropriately in a wide array of different circumstances, hold beliefs of the appropriate sort, and so forth.²⁸

Of course there may be good reasons to believe that the apparent functional profile of the virtues is misleading. And among these reasons may be ones that make no appeal to what I've called the "epistemic claim." Consider, as a case in point, so-called causal exclusion arguments so prominently discussed in the philosophy of mind.²⁹ As they are usually developed, versions of this type of argument purport to establish two conclusions: first, that the existence of causally efficacious *sui generis* supervenient entities of any sort implies either massive causal overdetermination, violations of the causal closure of the physical, or objectionable types of downward causation. And, second, these conclusions are sufficiently repugnant that we ought to accept that such putatively supervenient features are either epiphenomenal or type-identical with (or entirely explainable by) the physical features that ostensibly determine them.

Causal exclusion arguments arguably offer the most promising manner by which to develop Audi's suspicion without appeal to the epistemic claim.³⁰ Nonetheless – and this is to address the second concern raised earlier – I doubt that moral realists should be much enamored with them. Few realists, I suspect, would be enthusiastic about a view according to which both mental and moral features are type-identical with (or entirely explainable by) the physical facts that determine them; the problem being that it is difficult to see how, according to such a view, we could account for either the phenomenal aspects of mentality or the normative aspects of morality in a materialistically acceptable way.³¹ And few realists, I suspect, would be willing to believe that *both* mental and moral features are mere shadows of the subvenient physical realm. Audi, for one, *defends* the view that higher-order entities such as beliefs are causally efficacious.³² In saying this, I don't mean to claim that reductionism or epiphenomenalism is strictly inconsistent with moral realism. I mean only to suggest that the reasons for (affirming and alternatively, denying) causalism with respect to the virtues and mental features are so similar that it is doubtful we can affirm one and not the other. Accordingly, if one wants to affirm that mental features have the phenomenal features we ordinarily ascribe to them and are causally efficacious, then one should probably also embrace causalism. And, conversely, if one wishes to reject causalism, then one should probably be prepared to reject our ordinary understanding of the mental.

VI. *A loose thread and a suggestion*

I would like to close by tying up a loose thread and then making a suggestion about how to extend the present argument.

First, the loose thread: Earlier I said that my aim in the last section was to defend the following conditional claim: If there is a configuring cause of the fact that Sam is motivated appropriately, then it is the exercise of Sam's compassion. But, so far, I have not defended the claim that there is a configuring cause of the fact that Sam is motivated appropriately. Is there anything further to say in its defense?

In fact, I have at this point, little new to say in its defense. I don't think the claim should be rejected because there is something suspect with the notion of a configuring cause. To the contrary, if the preceding discussion has been on the mark, the concept appears to do important explanatory work insofar as it can help us to address problems such as accidentally reliable connections and deviant causal chains. Nor do I think the claim should be rejected because there is no causal explanation of the fact that Sam is motivated appropriately. For this latter claim to be sustained a version of moral epiphenomenalism would have to be true. But if what I've argued is sound, realists should no more be attracted to epiphenomenalism with respect to the virtues than to epiphenomenalism with respect to the propositional attitudes.

Still, I suspect that the argument I've developed, even if sound, may leave some of those sympathetic with causalism unsatisfied. The lack of satisfaction, I suspect, is rooted in the thought that a satisfactory causalist view should not merely defend the claim that moral facts causally explain other moral facts, but ought also to defend the thesis that moral facts causally explain non-moral facts. Although I do not see why it is incumbent upon a causalist who does not wish to argue for moral realism to argue for this latter claim – call it the “standard view” – one could view the argument in this paper as the thin wedge of an argument whose aim is to establish that some moral facts cause non-moral facts of certain kinds. The argument would attempt to establish that the causal efficacy of moral facts of certain kinds is general in character: Roughly, given that we have good reasons to believe that moral facts of certain kinds cause other moral facts, and that there are also explanatory contexts in which we cite moral facts of certain kinds as the cause of non-moral facts, then, in the absence of reasons to think otherwise, we should also believe that moral facts of certain kinds cause non-moral facts of certain types. According to this strategy, we argue *to* the standard thesis from the phenomenon of causal relations between moral facts. Of course, if we implement only this strategy, the standard view cannot be used to argue for moral realism. But this may be a price we should be willing to pay, for implementing this strategy may be the most promising avenue by which we can develop a plausible moral naturalism.³³

Department of Philosophy
Calvin College

NOTES

¹ See Harman 1977, chap. 1; Williams 1985, chap. 8; Gibbard 1990, chap. 6; Zimmerman 1985; Wright 1992 and 1995; and Blackburn 1995.

² That the “explanatory challenge” concerns the causal efficacy of moral facts becomes clear when one examines the realist responses to it. See, for example, Sturgeon 1985 and 1986; Brink 1989, chap. 7; McGinn 1997, chap. 2; and Shafer-Landau 2003, chap. 4.

³ See Brink 1989, chap. 7; Boyd 1988; Sturgeon 1985; Railton 1986; and Jackson 1998, chaps. 5 and 6.

⁴ The claim is made by Kim 1996, p. 130 among many others. In what follows, I shall not assume this claim is true.

⁵ See Dworkin 1996; McGinn 1997, chap. 2; Audi 1997 and 1993; and Thomson 1996.

⁶ See, e.g., Railton 1986; Sturgeon 1985; and Brink 1989, chap. 7.

⁷ See, e.g., Dupré 1993; Cartwright 1983 and 1999; Prades and Corbí 2000; Baker 1995; and Mellor 1995.

⁸ See Hale 1987, p. 94.

⁹ See Dworkin 1996, p. 119.

¹⁰ See Thomson 1996 and Audi 1993.

¹¹ See McGinn 1997, pp. 13–16.

¹² Sturgeon 1985 also suggests that virtues and vices are plausibly seen as being causally efficacious. One might view the causalist position I wish to defend as elaborating upon Sturgeon’s suggestion.

¹³ Cf. Dretske 1993.

¹⁴ Here and elsewhere I shall remain neutral concerning different accounts of what “informational content” consists in; I certainly don’t wish the example I employ to suggest that the informational content of human cognitive systems can be satisfactorily analyzed as symbolic representations processed by automated formal symbol-manipulation systems!

¹⁵ Marc Slors has pointed out to me that what I call a “configuring” cause bears a resemblance to what Fred Dretske calls a “structuring” cause (see Dretske 1993). Nonetheless, the two concepts are different: A structuring cause is the cause of a more or less stable condition that explains why a triggering cause A causes B. In the case I have described, the structuring cause of letters of certain types appearing on a given computer’s monitor is the fact that someone has programmed the computer in a certain way. The configuring cause of the fact that letters of a certain kind are produced on that computer’s monitor, by contrast, is the fact that certain features of the computer’s hardware and software are working in a certain way.

My preference for working with the concept of a configuring rather than a structuring cause lies (in part) in the desire to avoid committing the realist to the somewhat controversial claim that capacities or dispositions can themselves be causes. That said, if one preferred to employ the concept of a structuring cause, one could stipulate that a capacity of a system S is a structuring cause of a process in S only in those cases in which (i) no interference conditions are in place that block the exercise of this capacity and (ii) the capacity doesn’t malfunction in certain ways.

¹⁶ I use the term ‘capacity’ in the same fashion that Cartwright 1999 does. Capacities are “open-ended” insofar as they are not restricted to any single type of manifestation.

¹⁷ Here is one worry concerning this suggestion: How could it be informative to claim that what causally accounts for the fact that Sam is motivated appropriately over some stretch of time is the fact that he has manifested the virtue of compassion – a disposition, in part, to be motivated aright?

In the following two ways: First, it is important to note that there are interesting competing explanations for why Sam is motivated aright – explanations I will address in

sections IV and V. One such competing explanation is that there is a configuring cause of his being motivated appropriately, but it does not involve the exercise of a virtue. Another competing explanation is that there is no configuring cause of Sam's being motivated appropriately. One reason, then, that the putative causal explanation I am broaching is informative is because it has what van Fraassen 1980 calls a relevant "contrast class."

Second, it is helpful to keep in mind that the term "motivated appropriately" is being used as shorthand for a descriptively rich constellation of motivational states that are the product of a descriptively rich cluster of experiential inputs. To say that Sam is motivated appropriately is to say that, given his awareness of a variety of signs, Sam is appropriately motivated to listen carefully to what Margaret says in conversation, call upon her regularly for coffee, and so forth. And it is not trivial to say that the exercise of Sam's virtue causally accounts for the fact that this descriptively rich cluster of motivational states forms in response to a wide variety of different experiential inputs. For an elaboration of this basic idea, see Hempel 1965, p. 457ff.

¹⁸ Here I echo Railton 1995, p. 276. Perhaps it is worth explicitly noting that the type of explanation in question doesn't seem to be merely a normative explanation. The exercise of a virtue doesn't serve merely to *justify* these constancies in behavior. Nor does the explanation seem teleological in character. It is not as if the exercise of the interpretive and perceptual aspects of Sam's virtue somehow detects the compelling rightness of Sam's being motivated appropriately.

¹⁹ According to the canonical understanding, thick moral properties such as the virtues are an amalgam of both descriptive and normative elements. Accordingly, one way to defend the claim that virtues are causes is to argue that they are causally efficacious only in virtue of their descriptive elements. I assume, however, that those realists who reject the causalist view would not wish to object to this claim. As I understand it, their contention is that moral facts *qua* normative entities – that is, *qua* moral reasons or morally appropriate responses to reasons – are not causally efficacious.

²⁰ See Audi 1997 and 1993. Thomson 1996 has also formulated an objection to the causalist position that more or less mirrors Audi's. For ease of explication, I shall consider only Audi's version of the argument. What I say about Audi's view should apply to Thomson's position *mutatis mutandis*.

²¹ Audi 1997, pp. 118–19. See, also, Audi 1993, p. 62. Audi says here that police brutality and the like should be *construed* behaviorally, and that we can explain the revolt in terms of police brutality and the like. This suggests that Audi thinks that we can explain the revolt in purely behavioral terms. But since I doubt that there is a plausible explanation of the revolt that refers only to behavioral facts, I shall interpret Audi to claim that, although we can *identify* police brutality and the like behaviorally, we cannot *explain* the revolt in the absence of those propositional attitudes that constitute the brutality, deception, and so forth.

²² Audi 1997, p. 118.

²³ Cf. Audi 1993, p. 62.

²⁴ As both Elgin 1996 and Zagzebski 2001 point out, true beliefs can actually be an impediment to understanding a situation well. Sometimes a representation of a situation that is strictly speaking false, because oversimplified, allows us a better understanding of a situation than one that is true.

²⁵ For more on these matters, see Plantinga 1993 and Zagzebski 1996, sect. 4.

²⁶ For more on this theme, see Plantinga 1993; Zangwill 1996 and, especially, Wedgwood (unpublished).

²⁷ It is this fact that speaks against construing configuring cause explanations as a species of what Jackson and Pettit (1990) call "program explanation." Roughly put, program explanations appeal to higher-level features (such as temperature) to isolate lower-level

features (such as a certain kinds of molecular activities) that are doing the causal explanatory work. I have argued, however, that the shapelessness of the normative does not “program” for these lower level features, at least not under any description that is genuinely informative (i.e., not of the form “whatever natural features that determine the exercise of the virtue in question”).

²⁸ See Burge 1993, p. 112. See also Brink 1989, pp. 194–97.

²⁹ For a development of (and ultimately a rejection of) the argument as applied to the moral realm, see Shafer-Landau 2003, chap. 4. The classic formulation of the exclusionary argument is due to Kim 1993, chaps. 6, 14, and 17, and 1998.

³⁰ In one place, Audi says that if “moral properties are . . . not *candidates* to be epiphenomenal: roughly, since they are not causally dependent variables, there is no reason to lament if they are not causal variables at all; they are apparently not the right sort of property to be in either category” (Audi 1997, p. 122). I am not sure that I understand Audi’s thought here, however. Perhaps the idea is that, since the supervenience relation that characterizes the supervenience of the moral on the non-moral is not a species of causal-nomic dependence, it follows that moral facts are not candidates for being in causal relations. If that is what the claim is, then I see no reason to believe it. And, to my knowledge, Audi does not provide an argument for believing it is true.

³¹ Although see Jackson 1998, chaps. 5 and 6. For a response that develops the concerns raised here, see Majors 2005.

³² See Audi 1993.

³³ Thanks to Jason Baehr, Igor Duven, Chris Eberle, Phil Goggans, Steve Layman, Brad Majors, Luke Reinsma, Ron Rood, Marcelo Sabates, Michael Schwiegger, Russ Shafer-Landau, Mark Slors, René van Woudenberg, Pekka Väyrynen, my colleagues in the philosophy department at Calvin College, and an audience at the conference *Between Armchair and Experiment: Cross-Fertilization between Philosophy and Science*, Heeswijk, Netherlands, May, 2001, for their comments on an earlier draft of this essay.

REFERENCES

- Audi, R. (1997). “Ethical Naturalism and the Explanatory Power of Moral Concepts,” in his *Moral Knowledge and Ethical Character*. Oxford: Oxford University Press, pp. 112–128.
- Audi, R. (1993). “Mental Causation: Sustaining and Dynamic,” in Heil and Mele (eds.) (1993), pp. 53–74.
- Baker, L. R. (1995). *Explaining Attitudes*. Cambridge: Cambridge University Press.
- Blackburn, S. (1995). “The Flight to Reality,” in R. Hursthouse, *et al.* (eds.) *Virtues and Reasons*. Oxford: Oxford University Press, pp. 35–56.
- Boyd, R. (1988). “How to be a Moral Realist,” in Sayre-McCord (1988) (ed.), pp. 181–228.
- Brink, D. (1989). *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Burge, T. (1993). “Mind-Body Causation and Explanatory Practice,” in Heil and Mele (eds.) (1993), pp. 97–120.
- Cartwright, N. (1999). *The Dappled World*. Cambridge: Cambridge University Press.
- Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford: Oxford University Press.
- Copp, D. and D. Zimmerman, eds. (1985). *Morality, Reason and Truth*. Totowa, NJ: Rowman and Allanheld.
- Cordí, J. and J. Prades. (2000). *Minds, Causes and Mechanisms*. Oxford: Blackwell.
- Dretske, F. (1993). “Mental Events as Structuring Causes of Behaviour,” in Heil and Mele (eds.) (1993), pp. 121–36.

- Dupré, J. (1993). *The Disorder of Things*. Cambridge, MA: Harvard University Press.
- Dworkin, R. (1996). "Objectivity and Truth: You'd Better Believe It," *Philosophy and Public Affairs* 25, pp. 87–139.
- Elgin, C. (1996). *Considered Judgment*. Princeton: Princeton University Press.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press.
- Hale, R. (1987). *Abstract Entities*. Oxford: Blackwell.
- Harman, G. (1977). *The Nature of Morality*. Oxford: Oxford University Press.
- Harman, G. and J. J. Thomson. (1996). *Moral Relativism and Moral Objectivity*. Oxford: Blackwell.
- Heil, J. and A. Mele, eds. (1993). *Mental Causation*. Oxford: Oxford University Press.
- Hempel, C. (1965). *Aspects of Scientific Explanation*. New York: The Free Press.
- Jackson, F. (1998). *From Metaphysics to Ethics*. Oxford: Oxford University Press.
- Jackson, F. and P. Pettit. (1990). "Program Explanation: A General Perspective." *Analysis* 50, pp. 107–17.
- Kim, J. (1998). *Mind in a Physical World*. Cambridge, MA: MIT Press.
- Kim, J. (1996). *Philosophy of Mind*. Boulder, CO: Westview Press.
- Kim, J. (1993). *Supervenience and Mind*. Cambridge: Cambridge University Press.
- Majors, B. (2005). "Moral Discourse and Descriptive Properties." *The Philosophical Quarterly* 55, pp. 475–94.
- McGinn, C. (1997). *Ethics, Evil and Fiction*. Oxford: Oxford University Press.
- Mellor, D. H. (1995). *The Facts of Causation*. London: Routledge.
- Plantinga, A. (1993). *Warrant and Proper Function*. Oxford: Oxford University Press.
- Railton, P. (1985). "Moral Realism," *Philosophical Review* 95, pp. 163–207.
- Railton, P. (1995). "Subjective and Objective," *Ratio* 8, pp. 259–76.
- Sayre-McCord, G., ed. (1988). *Essays in Moral Realism*. Ithaca, NY: Cornell University Press.
- Shafer-Landau, R. (2003). *Moral Realism: A Defence*. Oxford: Oxford University Press.
- Sturgeon, N. (1985). "Moral Explanations," in Copp and Zimmerman (eds.) (1985), pp. 49–78.
- Sturgeon, N. (1986). "Harman on Moral Explanations of Natural Facts," *Southern Journal of Philosophy* 24, supplement, pp. 69–78.
- Thomson, J. J. (1996). "Epistemological Arguments for Moral Scepticism," in Harman and Thomson (1996), pp. 69–94.
- Van Fraassen, B. (1980). *The Scientific Image*. Oxford: Clarendon Press.
- Wedgwood, R. (unpublished). "Normative Explanations."
- Williams, B. (1985). *Ethics and the Limits of Philosophy*. Cambridge, MA: Harvard University Press.
- Wright, C. (1995). "Truth in Ethics," *Ratio* 8, pp. 209–26.
- Wright, C. (1992). *Truth and Objectivity*. Cambridge, MA: Harvard University Press.
- Zagzebski, L. (1996). *Virtues of the Mind*. Cambridge: Cambridge University Press.
- Zagzebski, L. (2001). "Recovering Understanding," in M. Steup (ed.), *Knowledge, Truth, and Duty*. Oxford: Oxford University Press, pp. 235–54.
- Zangwill, N. (1998). "Direction of Fit and Normative Functionalism," *Philosophical Studies* 91, pp. 173–203.
- Zimmerman, D. (1985). "Moral Realism and Explanatory Necessity," in Copp and Zimmerman (eds.) (1985), pp. 79–103.