

MARC LANGE

SALIENCE, SUPERVENIENCE, AND LAYER CAKES IN  
SELLARS'S SCIENTIFIC REALISM, MCDOWELL'S MORAL  
REALISM, AND THE PHILOSOPHY OF MIND

(Received in revised form 20 May 1998)

My story begins with the “layer-cake” conception of science. In Hempel’s words:

[I]t will be helpful to refer to the familiar rough distinction between two levels of scientific systematization: the level of *empirical generalization*, and the level of *theory formation*. ... [T]he former level ... is characterized by the search for laws (of universal or statistical form) which will establish connection among the directly observable aspects of the matter under study. [In] the second level ... research is aimed at comprehensive laws, in terms of hypothetical entities, which will account for the uniformities established on the first level (1965, p. 178).

From this layer-cake picture of scientific theorizing, there arises Hempel’s “Theoretician’s Dilemma”. Consider a law at the higher (i.e., “theoretical”) level: “All  $T_1$ ’s are  $T_2$ ’s”. When this theoretical law is supplemented by various “bridge” laws (such as “The  $T_1$ ’s are exactly the  $O_1$ ’s” and “The  $T_2$ ’s are exactly the  $O_2$ ’s”), it entails various laws relating kinds at the lower (i.e., observational) level, such as “All  $O_1$ ’s are  $O_2$ ’s”.<sup>1</sup> But then this law, the empirical counterpart of the theoretical law, generates from our observations ( $O_1$ ) all of the empirical predictions ( $O_2$ ) that the theoretical apparatus gives us. The theory, while convenient, seems in principle utterly dispensable for making empirical predictions. A rational reconstruction of scientific work should apparently declare the entire theoretical layer, with its discussion of posited unobservable entities, to be superfluous – contrary to scientific realism. As Hempel expresses the challenge:

Why should science resort to the assumption of hypothetical entities when it is interested in establishing predictive and explanatory connections among observ-



ables? Would it not be sufficient for the purpose, and much less extravagant at that, to search for a system of general laws mentioning only observables, and thus expressed in terms of the observational vocabulary alone? (1965, p. 179)

In a famous series of papers (1963a, 1963b, 1965, 1977), Sellars argues that theories are not dispensable for predictive purposes; rather, the layer-cake picture is mistaken. As I shall explain in section I, Sellars argues that contrary to the layer-cake picture, the empirical counterparts of theoretical laws cannot be ascertained *autonomously* – that is, without guidance from theoretical considerations. The kinds figuring in the empirical counterparts of the theoretical laws are salient only from the theoretical outlook. The empirical generalizations that *are* arrived at by working autonomously at the empirical level are not natural laws, even if they are true. The kinds figuring in these generalizations do not cut nature at its joints.

I argue in section II that the difficulty Sellars identifies for the layer-cake view of science is analogous to the difficulty that moral projectivism encounters according to McDowell (and Wiggins (1987)). Indeed, I have just expressed Sellars's point in McDowell's terms, such as a kind's "salience" from a certain "outlook" (1979, p. 345). Moral projectivism involves its own layer cake, with a lower layer of naturalistic kinds and an upper layer of moral kinds. McDowell (1978, 1979, 1981, 1985, 1987) suggests that although the moral kinds supervene on the naturalistic kinds, the naturalistic counterpart of a moral kind is not rendered salient from an autonomously naturalistic framework.

In section III, I explain that Sellars's argument is also the model for an argument in the philosophy of mind best approximated by Fodor (1974). Again, the issue is the relation between two layers: a lower layer of microphysical (or neurological) kinds and an upper layer of folk-psychological categories. One way of expressing a common view, it seems to me, is that the generalizations of folk psychology involve kinds that supervene on microphysical kinds but are not salient at the physical level, because of the diversity of the physical states that can realize a given folk-psychological kind. The microphysical counterpart of a folk-psychological generalization is not a law of physics; we cannot arrive at it by working autonomously

at the physical level. I examine this argument in the third section of this paper.

I have two motivations for displaying the analogies among these three arguments.<sup>2</sup> First, I wish to understand these arguments better. By juxtaposing them, I distinguish the elements essential to their (common) basic structure from those that are dispensable or even misleading. We can better evaluate various objections to these arguments by running them in parallel – by noticing, for example, that a difficulty van Fraassen raises against Sellars is analogous to a difficulty Blackburn raises against McDowell. My second motivation is to understand precisely what it would take for these arguments to succeed. All three depend upon the same presuppositions about laws of nature, natural kinds, and salience. In section IV, I explore briefly what natural laws and natural kinds would have to be for these arguments to go through.

In particular, the natural kinds would have to be the kinds figuring in the natural laws, and a natural law would have to be salient from some standpoint or other. This latter idea – that for any natural law, there must be some standpoint from which, on sufficient evidence, a correct inductive argument would justify belief in that law – is suggested by the traditional intuition that a law must be capable of deriving inductive support from its instances. This connection between salience and lawfulness would also account for the natural laws' failure to be closed under logical consequence, as well as their failure to supervene on the singular, local, non-nomic facts.

The three analogous arguments I discuss suggest one way for there to be real moral and folk-psychological kinds that are distinct from microphysical kinds, without requiring any spooky moral facts or mental stuff for the moral and mental laws to govern. Of course, the microphysical kind corresponding to a given moral or folk-psychological kind is wildly disjunctive, gerrymandered, or “fat”. (Furthermore, the moral or folk-psychological laws in which the given kind figures contain *ceteris-paribus* clauses, and the cases where “all other things are equal” are microphysically heterogeneous.) But this does not signal that the moral or folk-psychological kind is unreal, only that it is not a natural *microphysical* kind. The real moral and folk-psychological kinds are autonomous – governed by laws at their own layer that do not correspond to microphys-

ical laws – *precisely because* the members of a given moral or folk-psychological kind are microphysically heterogeneous.

# I

Sellars argues against two versions of the layer-cake picture. The first version is associated with phenomenalism; the lower-level laws connect kinds characterized in purely phenomenal terms, and the upper-level laws connect kinds characterized in terms of ordinary middle-sized objects. (This version of the layer-cake picture presumes that we do not directly observe ordinary middle-sized objects; it presumes that observation reports are made in purely phenomenal terms.<sup>3</sup>) On the second version of the layer cake, the lower-level laws relate kinds characterized in terms of ordinary middle-sized objects, and the upper-level laws relate kinds characterized in terms of posited unobservable entities. In opposing this version of the layer-cake view, Sellars is arguing against scientific anti-realism.

Sellars offers the same argument against both versions; it is directed against the layer-cake structure they have in common. In that structure, the bridge laws need not be biconditionals (such as “The  $T_1$ ’s are exactly the  $O_1$ ’s”), but they could be. Sellars considers such an ideal case; he concedes, for the sake of argument, that an upper-level kind ( $T_1$ ) can be demarcated in lower-level terms ( $O_1$ ) – i.e., that a kind specifiable in theoretical terms is, as a matter of physical necessity, coextensive with a kind specifiable in empirical terms. Sellars is thereby setting aside a notorious difficulty for the layer-cake view: that there may well be no bridge laws because it is doubtful that the requisite necessary or sufficient conditions for membership in a theoretical kind can be given in observational terms. Having put aside this difficulty, Sellars refers to certain empirical generalizations as the “observational counterparts” (1963a, p. 83; 1977, p. 320) of theoretical laws.

But here lies the trouble, according to Sellars: Although the evidence includes many instances and no violations of the observational counterparts of theoretical laws, these empirical generalizations are not *suggested* by the evidence. Consider, for example, the phenomenal counterpart of some familiar law linking material-object categories, such as “All lightning is followed by thunder.” Many

phenomenal predicated pick out from my perceptual history exactly the cases in which I have observed lightning. Some of these predicates are suggested by the record of my perceptual history; others make unmotivated bends and so are not rendered salient. When does my evidence justify my believing that “All  $O_1$ ’s are  $O_2$ ’s” is a law, where this generalization is compatible with my evidence, and where the  $O_1$ ’s in my perceptual history are exactly the cases in which I have observed lightning (and the  $O_2$ ’s thunder)? Only when the evidence justifies my selecting  $O_1$ ; that is, only when  $O_1$  is suggested by my perceptual history. To adopt a generalization whose antecedent makes unmotivated bends would be unmotivated – that is, unjustified. A hypothesis with a gerrymandered antecedent is not confirmed by its instances in my perceptual history. Sellars’s point is that the antecedent of the phenomenal counterpart of a material-object law will be gerrymandered from a phenomenal standpoint; my perceptual history, without theoretical considerations, will *not* make it salient.

Why is this? Consider a phenomenal property that *is* rendered salient by my perceptual history. If this property covers all my past observations of lightning but none of my other past observations, then (Sellars plausibly holds) it does so only because of incidental features of my history. How, for example, do my observations of lightning differ in phenomenal terms from my observations of various photographs of lightning (which, of course, are not invariably followed by my hearing thunder)? Perhaps, as it happens, all of the photographs of lightning I have observed have been in photography books, printed on glossy white paper, and so have been surrounded by glossy white sense-data. Perhaps it also happens that none of my observations of lightning has been through a window set in a wall covered by glossy white wallpaper or paint. Then the phenomenal regularity suggested by my perceptual history might well be (roughly speaking) that all lightning-shaped sense-data not surrounded by glossy white sense-data are followed by thunder-like sense-data. But it is mere happenstance that this regularity has obtained. Sellars writes:

The uniformities I find are bound up with the fact that my environment has included wallpaper of such and such a pattern, a squeaky chair, this stone fireplace, etc. etc. My having had *this* pattern of sense contents has usually eventuated

in my having had *that* sense content, because having *this* pattern of sense contents guarantees, for example, that I am ... looking at the fireplace (1963a, p. 82).

A phenomenal generalization that is suggested by my perceptual history has held only accidentally; even if it remains true throughout my future perceptual career, it expresses merely an accidental regularity, not a natural law, since it obtains only because life happens never to take me into certain environments. (Sellars calls such a regularity “essentially autobiographical” (1963a, p. 83), as distinct from a non-accidental – i.e., physically necessary – regularity in my perceptual history, which he calls “accidentally autobiographical.”) Since the phenomenal generalization that is confirmed by my perceptual history reflects accidental features of my life thus far, it will hold of my future perceptual career (and of others’ careers) only if the rest of my life (and their lives) happen also to possess those features. Therefore, the phenomenal generalization that is confirmed by my perceptual history typically makes inaccurate predictions regarding my future observations and others’ observations. Typically, each time I expand my evidence to include new observations that I or others have made, the phenomenal generalization that I have justly adopted so far is overturned. Sellars refers to the “instability” (1977, p. 314) of these generalizations: “they will exhibit an incorrigible variance with respect to fresh cases” (1977, p. 320).

Sellars’s point is not that there are no physically necessary lower-level generalizations, i.e., that there are no purely empirical counterparts of theoretical laws. Rather, his point is that I cannot justly arrive at physically necessary lower-level generalizations by inductively projecting regularities holding among my past observations unless I am guided in my projections by an upper-level outlook. The empirical counterparts of the kinds figuring in theoretical laws are rendered salient by our observations if and only if we allow theoretical considerations to play a role in our inductive reasoning. Hence, even if the theoretical layer adds nothing to the capacity of the lower-level generalizations to make correct empirical predictions, the theoretical layer is indispensable to our making those predictions, since it is crucial to our justly arriving at those lower-level generalizations. Sellars says:

[T]he very selection of the complex pattern of actual sense contents in our past experiences which are to serve as the antecedents of the generalizations in ques-

tion presuppose our common sense knowledge of ourselves as perceivers, of the specific physical environment in which we do our perceiving ... (1963a, p. 84).

Of course, this point is not restricted to the selection of antecedents; van Fraassen (1977, p. 337) mentions a good example – Van der Waals's gas law – in which the consequent is not salient on “autonomous inductive reasoning in the observation framework” (Sellars 1977, p. 319). Here we are not dealing with phenomenalism; in this example, the material-object generalizations constitute the empirical layer, and generalizations involving microphysical kinds constitute the theoretical layer. Recall Van der Waals' law: All gases conform to  $(P + a/V^2)(V - b) = RT$ , where  $P$  is pressure,  $V$  is volume per mole,  $T$  is temperature,  $R$  is the gas constant, and  $a$  and  $b$  are parameters characteristic of the particular chemical species. The consequent of this law does not “contain empirical concepts which it would be reasonable to construct and use in the absence of theoretical considerations” (Sellars, 1965, p. 194). From an autonomously material-object outlook, the concepts Van der Waals uses appear gerrymandered.

Let me elaborate this example more carefully, to bring out Sellars's point. Van der Waals justified belief in this law by employing a particular “inductive strategy”. This strategy seeks a  $P$ - $V$ - $T$  equation (an “equation of state”), applicable to any gas, that is sufficiently reliable for a certain range of uses. The strategy proceeds from a theoretical working assumption: that all gases are composed of small, ceaselessly moving particles attracting each other by short-range forces, whose collisions with the container walls are responsible for the gas's pressure, etc. An equation is a candidate for salience on this strategy only if its reliability could be explained through this working assumption, i.e., only if the equation can be derived by adding to the ideal gas law ( $PV = RT$ ) various correction factors that take into account certain further influences posited by this working assumption. Different equations can be derived by taking different influences to be so large that they cannot be neglected and by using different approximations for generating the factors correcting for these influences. Van der Waals's equation, for example, results from taking the only non-negligible influences to be the gas molecules' sizes and mutual attractions, and by making certain other approximations (e.g., that the density of the gas's surface layer is

inversely proportional to the gas's volume).<sup>4</sup> The working assumption behind an inductive strategy – the assumption constituting (in Sellars's words) its "standpoint" or "framework" – must suffice to ensure agreement on which hypothesis, if any, is rendered salient by our observations, i.e., agreement on what it would be for gases (in this example) to behave in the *same way* in unexamined cases as they do in the cases we have already examined. Once an equation is rendered salient, the inductive strategy requires us to regard its instances as inductively confirming its reliability, i.e., as supporting the equation's "projection" onto unexamined cases, actual and counterfactual alike.<sup>5</sup> On the particular inductive strategy I have specified, Van der Waals's equation is quickly rendered salient. This would not have happened without some such theoretical motivation to limit the candidates for salience. This is Sellars's key point: on different inductive strategies, seeking the same kind of equation but proceeding within different frameworks, the same evidence can render different hypotheses salient.

That is: Many different equations were compatible with the experimental observations that Van der Waals cited as justifying belief in his equation. Only in view of certain theoretical considerations does Van der Waals's equation count as what it would be for gases under unexamined conditions to behave in the *same way* as they did under the range of conditions already examined experimentally. In particular, the very same observations that, taken in the light of theoretical considerations, suggested van der Waals's equation instead rendered Regnault's equation ( $PV = 1 - A[(1-V)/V] + B[(1-V)/V]^2$ ) salient on an autonomously observational framework. Whether an equation is suggested by (i.e. is rendered salient by) the data – whether it expresses what it would take for unexamined gases to behave in the same way as examined gases – is a normative matter, not a mere psychological one; it is a matter of which equation we are justified in projecting. Regnault's and van der Waals's equations both predict that in the range of pressures for which there existed experimental evidence,  $PV$  decreases as  $P$  increases (at a fixed temperature). If we are proceeding in the autonomously empirical framework, then (as noted by Loeb, 1934, p. 146) these data justify our expecting this trend to continue indefinitely; we should judge this to be what it would take



for the PV curve to go on in the same way. Consequently, this is what Regnault's equation predicts, whereas van der Waals's equation predicts that PV will eventually begin to rise with increasing P. Thus, from the autonomously empirical perspective, van der Waals's equation takes an unmotivated bend at pressures beyond the range of the available experimental data, and so is not rendered salient by these data, despite its compatibility with them. In contrast, this "bend" can justly be anticipated in the theoretical framework: at sufficiently high pressures, the gas particles become so crowded together that V cannot fall fast enough to keep pace with increasing P, and so PV begins to rise with P. What seems like a change in the gas's behavior, from an autonomously empirical viewpoint, appears from the molecular perspective as the gas's going on in the same way.

Does Sellars's argument undermine a certain variety of scientific anti-realism? Ultimately, I think it does, but not precisely as Sellars elaborates his argument. To explain why, I must briefly trace van Fraassen's objections and Sellars's replies to them. We can then better understand the significance of Sellars's argument.

Van Fraassen (1980, p. 34) notes that we can use a theory to guide our inductive inferences even if we have no opinion regarding its truth. Of course, van Fraassen is right if he means merely that we can use a theory as a *source of ideas* without presupposing anything about the likelihood of its truth. But Sellars's point is *not* that we would have *failed to think* of certain empirical categories if we had not been inspired by some theory. This would be merely an unfortunate psychological debility on our part. Sellars's point concerns not the context of discovery but the context of justification: to justify belief in the empirical counterpart of a theoretical law, we must be guided by theoretical considerations. How can we be justified in accepting guidance from a theory unless we have considerable confidence in its truth?

Van Fraassen (1976, p. 337; 1980) responds that it suffices for us to have some confidence in the theory's empirical adequacy. Sellars replies (1963b, p. 123) that belief in the theory's truth is needed to *explain* its empirical success. Van Fraassen is not persuaded by this response, since he denies that explanation belongs to a rational reconstruction of science; the only goal of science, according to van

Fraassen, is empirical adequacy. Since Sellars intends to argue that we need to proceed from a *theoretical* outlook in order to arrive at an *empirically adequate* theory, he considerably diminishes the force of his argument by later invoking explanatory power over and above empirical adequacy as science's goal. If he was ultimately willing to invoke explanatory power, he could have done so right from the start, as his initial response to the Theoretician's Dilemma: the theoretical level is needed for science to explain empirical regularities.

Instead of Sellars's appeal to explanation, I would offer an argument that appeals to nothing outside the anti-realist's austere conception of the goal of science. Suppose that the Y is some kind of posited unobservable entity (e.g., the electron), and suppose that all of our past observations are just as though there were Y's. For the past empirical success of the Y theory to confirm its future empirical success (i.e., for us to project the Y theory's empirical adequacy), we must believe that its past empirical success may be no coincidence; that is, we must believe that it may be a natural law, not merely an accidental generalization, that all phenomena are as if there were Y's.<sup>6</sup> But it is no coincidence that all observable facts are as if there were Y's *only if there really are Y's*. (This is the contrapositive of the familiar claim that if the Y is unreal, then it would be a coincidence if all observable facts were just as if there were Y's.) So for us to be justified in allowing the Y theory to guide our inductive projections, we must believe that the Y may be real; it is not enough for us to believe merely that the Y theory may be empirically adequate.<sup>7</sup>

Sellars's main point depends on the idea of salience: that not every regularity instantiated by our past observations is supported inductively by them. The key to Sellars's argument is that theoretical considerations can have a role in determining which regularity is salient, i.e., which depicts unexamined phenomena as going on in the same way as the phenomena we have examined. I shall take up this fundamental point again.

## II

The kinds of ordinary middle-sized objects figuring in the lower-layer laws supervene on the microstructural kinds figuring in the

upper-layer laws. Conceding (for the sake of argument) that for each of these subvening (microstructural) kinds there is a physically necessarily coextensive supervening kind, Sellars argues that if we proceed autonomously from a supervening outlook, the evidence will not make such a supervening kind salient. Having observed some cases of this kind, we will not regard all other cases of this kind as *like* the examined cases because what they have in common is gerrymandered as seen from the supervening perspective. While there is a supervening property that picks out exactly these cases, these cases belong together only from the subvening viewpoint.

Now switch “supervening” and “subvening” in this description of Sellars’s position: let the upper-layer kinds supervene on the lower-layer kinds, and consider autonomous inductive reasoning within the subvening framework. Take the supervening kinds to be moral categories and the subvening kinds to be naturalistic categories. You have arrived at McDowell’s position: He is rejecting a layer-cake picture in ethics by arguing, just as Sellars does, that only from an upper-layer perspective – in McDowell’s case, a moral standpoint – do the members of a given real upper-layer (moral) kind belong together.<sup>8</sup> Suppose we are taught a new evaluative category – e.g., courage – by being shown some examples of courage. Suppose we are then asked which further cases are *like* these examples, and so are also instances of courage. We will give wrong answers, McDowell contends, if we consider these examples from an autonomously naturalistic outlook, because from that perspective, courage is gerrymandered. If some naturalistic feature that the examples have in common is salient from an autonomously naturalistic outlook, the fact that this feature is shared by all of the examples of courage that we have been given is an essentially autobiographical regularity; had we been given another example of courage, it might well not have possessed this feature. There is a naturalistic kind that non-accidentally covers exactly the actual instances of courage, since it covers exactly the physically possible instances of courage: Specify naturalistically every physically possible situation (A, B, C, ...); identify naturalistically the instances of courage ( $\alpha$ ,  $\beta$ ,  $\gamma$ , ...) in A, the instances of courage ( $\aleph$ ,  $\beth$ ,  $\beth$ , ...) in B, etc.; and then “x instantiates courage” is non-accidentally coextensive with “x is ( $\alpha$

or  $\beta$  or  $\gamma$  or ...) in A or ( $\aleph$ , or  $\beth$  or  $\beth$  or ...) in B or ...".<sup>9</sup> But the resemblance among these cases can be appreciated only from the moral framework. McDowell writes:

Understanding why just those things belong together may essentially require understanding the supervening term (1981, p. 145).

We do not fully understand a virtuous person's actions – we do not see the consistency in them – unless we [have] a grasp of his conception of how to live (1979, p. 346).

Our inability to master a morally evaluative concept without grasping a certain "conception of how to live" might seem like a psychological idiosyncrasy or learning disability on our part, possessing no philosophical significance. Indeed, replying to McDowell, Blackburn has made an objection along these lines:

Let us suppose for a moment that some group of human beings does share a genuine tendency to some reaction in the face of some perceived properties or kinds of thing. Surely it need not surprise us *at all* that they should know of no description of what unifies the class of objects eliciting that reaction except of course the fact that it does so. We are complicated beings, and understand our own reactions only poorly. Now suppose the outsider, who fails either to share or to understand the reactive tendency, cannot perceive any such unifying feature either. Then he will be at a loss to extend the associated term to new cases, and there will be no method of teaching him how to do so. To take a very plausible example, it is notoriously difficult or impossible to circumscribe exactly all those things which a member of our culture finds comic. Any description is likely to have a partial and disjunctive air which would make it a poor guide to someone who does not share our sense of humor, if he is trying to predict what we will and will not find funny .... Do we really support a realist theory of the comic by pointing out the complexity and shapeless nature of the class of things we laugh at? On the contrary, there is no reason to expect our reactions to the world simply to fall into patterns which we or anyone else can describe (1981, p. 167).

Blackburn's reply to McDowell is like one of Van Fraassen's replies to Sellars: a moral or theoretical framework can be heuristically valuable in the context of discovery without playing any role in the context of justification; in view of our cognitive limitations, we may need to use that framework in order to think up a given hypothesis or to wrap our minds around the feature unifying the cases that would provoke a given response from us, but despite the framework's utility as a pump for our intuitions or reactions, it

can be omitted from a reconstruction of our *reasoning*. This version of Blackburn's objection<sup>10</sup> fails to recognize that McDowell's point concerns the context of justification. Sellars is concerned with how a prediction regarding an unexamined case is justified inductively by evidence that agrees with a general hypothesis entailing that prediction. McDowell's concern is similar: how a moral evaluation (e.g., "courageous") of some case under dispute (an "unexamined case") is justified by noting other things granted as courageous (the "examined cases") and then maintaining that the disputed case is *like* them. Both sorts of reasoning, scientific and casuistic, involve projecting some regularity onto unexamined cases, where from among the various regularities instantiated by the examined cases, we are justified in projecting only a regularity according to which the unexamined cases are *like* the examined cases. McDowell's point is that the examples granted as courageous will be alike in many respects; the salient respects determine which things we are justified in taking to be like the examples granted as courageous; and for the salient respects, the respects we are *justified* in seizing upon, to be those that lead non-accidentally to correct moral evaluations, we must already be approaching the examples from the moral standpoint:

To drop the primary-quality model [i.e., the layer-cake model – ML] in this case is to give up the idea that fearfulness itself, were it real, would need to be intelligible from a standpoint independent of the propensity to fear; the same must go for the relations of rational consequentiality in which fearfulness stands to more straightforward properties of things. Explanations of fear of the sort I envisage would not only establish, from a different standpoint, that some of its objects are really fearful, but also make plain, case by case, what it is about them that makes them so ... (1985, pp. 120f.).

A "relation of rational consequentiality" obtains between courageousness and various "more straightforward" (read; naturalistic) properties when something's possessing this combination of naturalistic properties is what "makes" it courageous and (here's the "rational" part) we can cite its possession of this combination of naturalistic properties to justify evaluating it as courageous.<sup>11</sup> But this justification goes through – we should recognize something's possession of these naturalistic properties as making it courageous – only if we recognize as similar the various combinations of nat-

uralistic properties that would suffice to make a thing courageous. We are not prepared to grant the similarity of all these combinations – we should consider them as forming an arbitrary, gerrymandered collection – unless we proceed from a moral standpoint, in which case there will be salient similarities.<sup>12</sup>

We might express these similarities either in explicitly evaluative terms – “All actions in which someone faces danger bravely, without flinching, are (*ceteris paribus*) courageous” – or by reference to particular examples – “An action is courageous if it is like rushing into a burning building to rescue the children trapped inside (*ceteris paribus*).” These moral generalizations are analogous to the natural laws in Sellars’ argument. Admittedly, these moral generalizations are not laws of nature, at least as natural laws are usually understood, but are intended to express explicitly certain norms governing the proper application of certain evaluative terms. Still, the analogy is appropriate, since we justify such a generalization by an analogue of induction: by showing how the generalization results from taking our intuitions regarding the evaluative term’s application to certain examples and projecting them onto “unexamined cases” so that its extension goes on in the same way. This is the sort of casuistic argument that appears in the Socratic dialogues with regard to such terms as “piety”. The analogy between McDowell and Sellars here is closer than you might think: in several papers, Sellars understands induction as making explicit, in the form of law-statements, the norms implicitly governing our use of various expressions.<sup>13</sup> To keep the analogy before us, let me refer to the generalizations on the moral side as “moral laws”.<sup>14</sup>

The two moral laws I have given (concerning courage) are typical in containing *ceteris-paribus* clauses. These clauses are needed because, for instance, someone rushing into a burning building to rescue the children trapped inside is not necessarily exhibiting courage if she is certain that she will not be hurt (because she is wearing a special suit and such rescues have become routine). The inclusion of a *ceteris-paribus* clause in a moral law does not trivialize that law, i.e., does not turn it into “An action is courageous if it is like rushing into a burning building to rescue the children trapped inside when doing so is courageous.” The *ceteris-paribus* clause is no different from any other predicate; it derives its content from a shared under-

standing of its range of application. That our knowledge of how to use the *ceteris-paribus* clause may be difficult to express explicitly does not suggest that the *ceteris-paribus* clause empties the moral law of content. As Kant and Wittgenstein have emphasized, the statement of a rule makes explicit how to apply a given concept only with the help of implicit norms governing the application of the concepts figuring in that statement; at some point, knowledge-how (or the faculty of judgment) must take over from knowledge-that. The *ceteris-paribus* clauses in moral laws are no different in this regard from those in scientific laws.<sup>15</sup> Consider Galileo's law of falling bodies: "All bodies falling to Earth accelerate at  $9.8 \text{ m/s}^2$ , in the absence of disturbing factors." Although this claim contains a *ceteris-paribus* clause, it is not empty, as it would be if it meant "All bodies falling to Earth accelerate at  $9.8 \text{ m/s}^2$ , except when they don't." Part of understanding the law of falling bodies is grasping an implicit norm governing its range of appropriate application, i.e., governing what qualify as "disturbing factors": another falling body colliding with the given falling body, air resistance, an electric field (where the falling body is electrically charged), a parachute slowing the fall, and so on. That Galileo's law is not made trivially true by its *ceteris-paribus* clause is evident from the fact that this claim was tested and found to be false; for instance, it does not hold at the poles, where gravity is stronger (because of the Earth's oblateness). Being at the North Pole is not what was originally meant by a "disturbing factor."<sup>16</sup>

Supervenience guarantees that there is some physical predicate<sup>17</sup> such that the range of physically possible cases to which it applies coincides with the range of physically possible cases that are courageous. But just as Sellars argues that the empirical counterpart of a theoretical law is typically not salient except from a theoretical viewpoint, so McDowell contends that the naturalistic counterpart of a moral law is typically not salient except from a moral viewpoint. From an autonomously naturalistic outlook, the cases covered by (e.g.) the *ceteris-paribus* clause in a moral law are not recognized as belonging together – this collection is gerrymandered – and so this moral law cannot be inductively inferred from cases. From some instances of the moral law, we might be justified in accepting some moral generalization with a *ceteris-paribus* clause, but the cases we

regard this clause as covering would typically not coincide with the cases covered by the *ceteris-paribus* clause in the genuine moral law unless we were working from the moral outlook.

Can we work within the moral outlook without believing moral properties *real*? This is analogous to a question van Fraassen asks Sellars: in justifying an empirical prediction, can we use a theory that we do not believe probably true, so long as we believe it probably empirically adequate? Surely, we can so use a moral theory to predict when someone else will apply a moral predicate (perhaps someone belonging to a culture with unfamiliar moral categories). We can use our understanding of their moral outlook to guide our inductive inferences – to influence which regularities in their past behavior are salient – even if we reject their moral outlook. That our justified inductive inferences yield inaccurate predictions of their evaluations, unless we proceed in their moral framework, does not demonstrate that the best explanation of their making some evaluation involves the reality of the moral properties they ascribe. Analogously, we might correctly postdict the reasoning of a past scientist only if we think about the world “in her terms”, but this fact does not suggest that her beliefs are true. The best explanation of our heuristic’s predictive success may simply be that it corresponds to her attitudes and beliefs.<sup>18</sup>

Once again, supervenience requires that for each evaluative predicate, there be some naturalistic predicate whose range in each physically possible world coincides with the evaluative predicate’s range. But this does not make the evaluative property “real” so far as I understand what’s supposed to be at stake in its “reality”. The naturalistic predicate that is non-accidentally coextensive with the evaluative predicate need not pick out a natural kind, just as the *emeruby* is not a kind of mineral even though the set of *emerubies* is non-accidentally the set of things that are emeralds or rubies.<sup>19</sup>

The argument for moral realism from considerations of salience seems sometimes to be thought of as proceeding in this way: We must believe in moral properties because we are disposed to make a certain response – to exhibit a certain moral attitude – towards exactly the objects in a certain set, and we lack a naturalistic criterion for membership in this set, so we must be detecting directly



that an object merits the response rather than ascertaining that the object possesses certain naturalistic properties and then responding to this information. Even if there is a purely naturalistic predicate that is non-accidentally coextensive with a moral predicate, our moral evaluations cannot be exhibiting our attitude towards that naturalistic property, since our recognizing this naturalistic property would have to be logically prior to our adopting some attitude towards it, whereas according to McDowell's argument from salience, our recognizing this naturalistic property is not logically prior to our making these moral evaluations.

But this argument for moral realism fails: talk of our adopting a certain attitude towards some naturalistic *property* is misleading. Our evaluations are of individual things; they need not be of some property that we must have already recognized as such. Here Blackburn's remark applies; our responsive dispositions can simply outrun our current capacity to express conceptually the range of circumstances that engage those dispositions. We may currently be unable to say what is common to all the things we would laugh at, except that they so move us. This failure is compatible with our responding, in each case, only to naturalistic properties. That we lack a naturalistic criterion for courage shows that we are not (always) *inferring* "x is courageous" from purely naturalistic claims, but fails to show that we are not *responding* to purely naturalistic properties when we are disposed to make this evaluation.

I added a final move to Sellars's argument: we can take the Y theory's empirical success in all examined cases as confirming its empirical correctness in all unexamined cases only if we believe it may be no accident that the Y theory has been empirically adequate, and so only if we believe that the Y may be real. An analogous addition to McDowell's argument is unhelpful, because of a crucial disanalogy between Sellars's argument and McDowell's. Sellars's argument is addressed to the scientific anti-realist, who acknowledges that we believe our observation reports of ordinary, middle-sized objects to describe the world and to be amenable to justification. My addition to Sellars's argument exploits a theory's role in justifying accurate predictions of observable facts. But McDowell's argument is addressed to the moral non-cognitivist; when McDowell argues that a moral theory is used to justify certain evaluations,

the non-cognitivist is prepared to grant neither that these evaluations purport to describe the world nor that they are amenable to justification.<sup>20</sup> So McDowell cannot presume a moral theory's role in justifying our moral evaluations, whereas the scientific realist can presume the Y theory's role in justifying our empirical predictions.

I think McDowell is best seen as arguing that moral evaluations purport to describe the world and are amenable to justification, and as so arguing by suggesting that the structure of moral thinking is analogous to the structure of science as depicted by Sellars's argument. Moral "laws" are no more trivialized by their *ceteris-paribus* clauses than are scientific laws, and they are no more amenable to justification outside the moral outlook than certain empirical generalizations in science are salient outside a theoretical perspective. The features of moral thinking that McDowell emphasizes, far from making the reality of moral categories suspect, are just the features needed to render morality more like science as Sellars reveals it to be.

McDowell says that arguments for or against moral realism concern "how ethics . . . relate[s] to the scientifically useful truth about the world and our dealings with it" (1987, p. 12) – how ethics is best "placed" (*Ibid.*) into what Blackburn calls "a working relationship with the scientific world view" (1993, p. 372). On the contrary, I see McDowell's argument as suggesting that the right way to defend the reality of moral properties is not to endeavor to fit them into the ontology associated with some familiar science – is not to argue, for example, that moral properties figure in scientific explanations of someone's behaviors. As I see it, McDowell's argument instead suggests that we compare moral thinking to familiar scientific thinking and examine whether they are structured similarly – in particular, whether there is on the moral side a distinction that performs a similar role to the distinction on the scientific side between real and artificial categories. (Arguments such as Sellars's begin to reveal the role that this distinction plays in science.) McDowell's argument suggests that there is such distinction on the moral side and, more fundamentally, that a real *moral* category need not play any role in a familiar science, since it is a real *moral* kind; it is salient solely from a moral viewpoint. Because there are distinct moral and naturalistic laws, there are distinct real moral and naturalistic kinds; if the kinds

that are salient from the moral outlook were not gerrymandered from a microphysical perspective, then nothing would be gained from occupying the moral perspective that is not already available through autonomous naturalistic inductive reasoning.

### III

This conception of moral kinds and naturalistic kinds as carving up the world differently, yet playing analogous roles (as “real” kinds) in separate disciplines, derives support from another argument of which Sellars’s is the prototype. This argument, in the philosophy of mind, seems to me the correct response to certain common thoughts about “folk psychology” (the sort of commonsense psychology that proceeds in terms of intentional states).

These thoughts concern the status of certain folk-psychological generalizations, such as “If someone believes that *p*, and believes that if *p* then *q*, then (*ceteris paribus*) she believes that *q*” and “If someone desires that *p*, believes that if *q* then *p*, and is able to bring it about that *q*, then (*ceteris paribus*) she brings it about that *q*.” Notice the *ceteris-paribus* clauses. For instance, someone whose other desires conflict with her desire that *p*, or who believes that *r* would be a better means than *q* for bringing it about that *p*, will not satisfy the *ceteris-paribus* clause in the second generalization. Someone who is confused or distracted, or cognitively impaired (though injury or illness etc.), will not satisfy the *ceteris-paribus* clause in any of these generalizations.

Since “cognitive impairment” is a disturbing factor, many philosophers have held that these *ceteris-paribus* clauses cannot be spelled out in terms of folk-psychological states, and so that there are no laws of folk psychology. For instance, Schiffer writes:

Will *every* nomologically possible physical defeater of these physical mechanisms itself realize a *psychological state*, such as confusion, irrationality, or distraction, that could take its place in a wholly *psychological* true completion of [a putative psychological law]? ... [N]o is the most plausible answer. ... [B]rain-injured people of the kind observed by Oliver Sachs appear to give empirical evidence that there are breakdowns in normal cognitive processes that cannot be accounted for in psychological terms (1991, p. 4).

(I return to this argument.) Suppose we cashed out the *ceteris-paribus* clause in physical terms. Each intentional state is realizable, consistent with the laws of physics, in an indefinitely large number of physical states. Hence, the physical cashing-out of “so long as there is no cognitive impairment” would have to be very long; it would have to list each physically possible kind of mental hardware and what it takes for that hardware to be working properly. But this clause would then, as Teller says, be too “fat” (1984, pp. 58f.) to appear in a statement of natural law; in an important sense, the various entries on the list have nothing in common (*Ibid.*). Millikan puts this common thought more fully:

Characteristically, the same function could, at least in principle, be performed by many differently constituted items. But . . . then the supporting conditions required for them to effect this function must differ as well. Brain cells performing the division algorithm require oxygen whereas computer chips require electric currents, and so forth . . . . The result is that there are no *ceteris paribus* laws governing all items having a certain function. For *ceteris paribus* conditions are unspecified conditions that must remain the same from case to case for the law to hold, whereas here the necessary conditions would have, precisely, to vary from case to case. A “law” applying to all such cases could say no more than that the items falling under the law could be made, by adding different circumstances tailored specifically to each case, to perform the function. But surely anything can be made to effect anything if one adds the right intervening media, if one adds enough special enough circumstances. So any such “law” would be empty (1993, p. 226f.).

According to this argument, the generalizations of folk psychology, with their non-psychological *ceteris-paribus* clauses, are not psychophysical laws. The same considerations of multiple realizability apply to replacing each intentional-state term in a folk-psychological generalization with a physical expression; the physical counterpart of “desiring that *p*”, for instance, would have to be a very long disjunction of disparate possible physical realizations – the realization would (in Millikan’s phrase) “vary from case to case.” It is therefore commonly held that there are no physical laws that are counterparts of the folk-psychological generalizations. By putting the point in terms of “counterparts”, I suggest the analogy to Sellars’s critique of the layer-cake picture. But when the ingredients in the above line of thought are assembled into an argument analogous to Sellars’s, the conclusion is that, contrary to Millikan,

there *are* folk-psychological laws. The ceteris-paribus conditions in these laws, while unrecognizable *from the lower-layer (microphysical standpoint)* as “remain[ing] the same from case to case”, are so recognizable *from the upper-layer (folk-psychological) standpoint*. Again, Sellars’s key point is that whether a category qualifies as gerrymandered depends on the framework.

I agree with Teller and Millikan that the physical counterparts of folk-psychological kinds are inadmissible into laws of physics because (as I would put it) they are not salient (they do not “remain the same from case to case”) from an autonomously microphysical perspective. If we note some instances of a folk-psychological generalization, considering those cases from a purely microphysical standpoint, then the generalization we are justified in adopting (involving kinds that “remain the same from case to case”) is essentially autobiographical;<sup>21</sup> that it holds of all the cases we have examined (or even all the actual cases) is an accident, a consequence of the restricted range of physical realizations we happen to have examined (or that happen to exist). Because the realizations in different hardware of the same intentional state are microstructurally so diverse, the particular realizations we have encountered (e.g., in human beings) will not allow us justly to anticipate the realizations we have never encountered (e.g., in space aliens), if we proceed within an autonomously microstructural framework. The long expression in microphysical terms that is *non-accidentally* coextensive with “desiring that p” (or with the ceteris-paribus clause in the folk-psychological generalization) does not appear in the hypothesis that is confirmed from an autonomously microphysical viewpoint. This category is salient – we will justly arrive at the physical counterpart of the folk-psychological generalization – only if we are guided by a theoretical (i.e., folk-psychological) outlook. Only from this outlook are all cases of “desiring that p” (or all cases that conform to the given ceteris-paribus clause) alike.

A folk-psychological generalization’s dependence on a theoretical perspective, far from impugning its lawfulness, makes it like a law of physics, as we saw in section I.<sup>22</sup> On this view, laws of intentional psychology bear much the same relation to laws of physics as moral laws bear to laws of physics: kinds that appear in the theor-

etical laws supervene on microphysical kinds, but a microphysical kind that is non-accidentally coextensive with one of these theoretical kinds is gerrymandered from an autonomously microphysical standpoint. There are distinct theoretical laws not because those theoretical kinds have no microphysical counterparts, but because those microphysical counterparts cannot be justly arrived at except via theoretical considerations.

Of course, these folk-psychological laws contain *ceteris-paribus* clauses. A cashing-out of one of these clauses in microphysical terms would produce a gerrymandered list. But why believe that these clauses need to be cashed out at all? Just as we understand (if we have learned basic physics aright) what qualifies as a “disturbing factor” in Galileo’s law of falling bodies, so we understand what qualifies as a disturbing factor in a folk-psychological generalization: not only confusion and distraction, but also illness, injury, and so on (as well as various other factors, depending on the generalization).<sup>23</sup> Schiffer worries (in the passage I quoted earlier) that injury (or the proper functioning of the nervous system) “cannot be accounted for in psychological terms.” If by “accounting for” an injury, Schiffer means giving a scientific explanation of it, then he is perhaps right, but this is irrelevant to the nomic status of folk-psychological generalizations. The laws of geometrical optics do not explain why mirrors reflect or why water refracts light; all the same, they govern those phenomena.<sup>24</sup> Perhaps by “accounting for” an injury, Schiffer means describing it “anatomically”. Once again, why is that important? Schiffer thinks that a folk-psychological generalization is empty unless we can express its *ceteris-paribus* clause more explicitly. But I see no reason why this must be so. Admittedly, for one of these clauses, we “would be hard pressed to say what the ‘other things’ are and what it is for them to be ‘equal’ ” (Schiffer, 1991, p. 2), if this means proceeding in microphysical terms and giving a complete list without resorting to “and so on.” But I have argued that this inability is intimately related to the reason why there are distinct folk-psychological laws in the first place. That the defeaters for different microphysical realizations are microphysically heterogeneous (i.e., any injury to one kind of organism, with its particular microphysical realization of a given intentional state, is anatomically different from

an injury to another kind of organism, with its distinct microphysical realization of that intentional state) does not entail that these defeaters are heterogeneous from a folk-psychological viewpoint; Sellars's argument turns precisely on the fact that what is salient from one standpoint is gerrymandered from another. And even without recourse to such an explicit list, we know how to recognize whether certain factors qualify as "disturbing", just as in using Galileo's law, we do not appeal to some complete list of disturbing factors.

Davidson, of course, is famous for disparaging *ceteris-paribus* clauses. He writes:

It is an error to compare a truism like "If a man wants to eat an acorn omelette, then he generally will if the opportunity exists and no other desire overrides" with a law that says how fast a body will fall in a vacuum. It is an error, because in the latter case, but not the former, we can tell in advance whether the condition holds ... (1990, p. 233).

The two cases seems entirely similar to me. One way of confirming that the fall is through a vacuum, that the body is not influenced by a strong electric field, etc., is to watch it fall. (This is mere confirmation, not conclusive proof, since the body would fall at  $9.8 \text{ m/s}^2$  if various disturbing factors existed but precisely counterbalanced each other.) Alternatively, we could draw upon our past observations and our knowledge of other laws (e.g., that strong electric fields exist between the plates of charged capacitors), and thereby confirm in advance whether (e.g.) the fall passes through a strong electric field. Likewise, to confirm that the man has no overriding desires and is cognitively intact, we might observe whether he eats the omelette under certain circumstances. Alternatively, since every folk-psychological law contains a *ceteris-paribus* clause requiring that the agent be cognitively intact, we could draw upon our knowledge of these other laws (and our past observations of the man's behavior) to confirm in advance whether he is cognitively intact. In the same way, we can use our knowledge of his past actions and relevant psychological laws to confirm in advance whether there are overriding conflicting desires.<sup>25</sup>

Drawing the analogy to Sellars's argument, I have referred to a moral or mental kind as having a microphysical "counterpart" with which it is non-accidentally coextensive. If there are moral and

mental laws along with laws of physics, then I must specify that by “non-accidentally” coextensive, I mean coextensive in all possible worlds with the same laws *of physics* as the actual world, since these are the laws governing the subvening (microphysical) kind. Notice, though, that the moral laws also hold in certain possible worlds with different laws of physics from the actual world’s – e.g., in a world in which the electron has twice the mass it possesses in the actual world. A circumstance specified in microphysical terms (e.g., there is an electron here, another there, and so on) that in the actual world would instantiate courage fails to do so in certain other possible worlds governed by different laws of physics but where the moral laws of the actual world still hold. For example, a microphysical circumstance that is dangerous in the actual world is not dangerous under certain other laws of physics. Consider the naturalistic kind that “non-accidentally coincides” with the actual instances of courage: “ $(\alpha \text{ or } \beta \text{ or } \gamma \text{ or } \dots)$  in A or  $(\aleph \text{ or } \beth \text{ or } \beth_1 \text{ or } \dots)$  in B or  $\dots$ ,” where A, B, C  $\dots$  consists of every possible situation consistent with the actual laws of physics. That something belongs to this kind exactly when it instantiates courage is not simultaneously a law of physics and a law of morals. More generally: For any physical property P, “All P’s instantiate courage” fails to be simultaneously morally necessary and physically necessary – and the same goes for “All instantiations of courage are P’s.” To be a law of morals, it would have to hold in all morally possible worlds, but these include some physically impossible worlds, including some where a microphysical circumstance that is dangerous in our world is not dangerous. If the generalization says that someone is courageous if she acts in a certain way under that circumstance, this is not so in certain morally possible worlds. If the generalization says that any instance of courage has some microphysical property P, then even if all physically possible instances of courage possess P, some instance of courage from a possible world with different laws of physics from the actual world’s will lack P. Likewise, there are no psychophysical laws in the sense of generalizations lawful in both physics and folk psychology. Just as the laws of statistical mechanics are “autonomous” in that they do not depend on what kinds of petty independent influences – what sort of “random” forces – are responsible for molecular motions, so the laws of folk psychology are auton-



ous in that they do not depend on details of the hardware, i.e., on what laws of physics govern the hardware realizing intentional states.<sup>26</sup>

In other words, the fact that all of the actual realizations of some intentional state are  $\alpha$  in hardware A or  $\beta$  in hardware B or... – where A, B... are all the hardware possible under the laws of physics – is an accidental generalization from the viewpoint of folk psychology. From that outlook, it just happens that all of the physical realizations of some intentional state are drawn from a certain limited range (namely, the range consistent with the actual laws of physics); a necessity in physics is an accident in psychology. To call folk psychology a “special science” (like “Earth science”), as if it were an application of physics to a certain restricted domain, is misleading; it suggests that necessities of folk psychology are accidents of physics while necessities of physics remain necessities in folk psychology, whereas in fact, necessities of physics are accidents in folk-psychology.<sup>27</sup>

#### IV

Of course, I have not made a detailed case for this conception of ethics and folk psychology. My purpose has been to lay bare the structure of Sellars’s argument and the analogous arguments concerning these other domains. It is then easier to see the presuppositions on which they all depend. I shall conclude by briefly examining the presupposition that serves as the fulcrum of this argument-scheme: the link between salience and natural law. These arguments presume that a given generalization states a law only if its instances can confirm it inductively and so only if, from among all of the rival generalizations compatible with those instances, they render that generalization salient from some “outlook.” For example, the microphysical counterpart of a folk-psychological law is not a law of physics because it is gerrymandered from a microphysical perspective; it cannot be arrived at by the inductive projection, from a microphysical perspective, of some of its instances, because from that perspective, the generalization does not portray unexamined cases as going on in the same way as examined cases. The question is: What is lawhood such that it is so linked to salience?

The distinction between natural laws and accidental generalizations is a venerable topic, and I cannot address it fully here. My specific concern is the constraint that lawhood's connection to salience imposes on any account of law – a constraint that perhaps cannot be met without sacrificing other strong intuitions. After all, why should it be logically impossible for a natural law to be gerrymandered? Couldn't there be a natural law that is not justly ascertainable by us? What's to stop nature from being so perverse? And since a generalization's salience ultimately depends on our justificatory practices, the natural laws would have to be essentially related to us if salience were necessarily linked to lawhood. Doesn't this overestimate our importance in the cosmos? Would the laws of nature change if our justificatory practices change? And what about counterfactuals? We surely cannot countenance the counterfactual: Had our justificatory practices been different, the laws of nature would have been different.

But perhaps the prospects for a link between salience and lawhood are not so bleak. For one thing, "natural law" might be an indexical expression referring rigidly to *our actual present* justificatory practices. In that case, there would be no pressure to acquiesce either to the laws' changing when our justificatory practices change or to the above counterfactual, any more than (to use Lincoln's example) we accept "Had a tail been called a 'leg', a calf would have had five legs." For another thing, lawhood's connection to salience would account for certain phenomena that are otherwise difficult to save. Let me briefly mention three.

First: Suppose the universe to be governed by Newton's laws. Consider the generalization: "Any two point masses attract each other with a force proportional to the product of their masses and inversely proportional to the square of their separation *unless* they are precisely 1.234 grams and 5.678 grams, in which case they attract each other with a force of 5 Newtons no matter how far apart they are." If there happens never to be point bodies of precisely those masses, then this "bent" generalization and Newton's gravitational-force law are both true, but only the latter is a law; only it correctly describes how two such bodies would have behaved, had they existed. Were lawfulness linked to salience, then the bent generalization's failure to qualify as a law would be explained by its

failure to be made salient by its instances; it does not portray point masses of 1.234 and 5.678 grams as behaving in the same way as bodies having other masses. Just as this generalization's instances do not suggest its prediction regarding unexamined point masses of 1.234 and 5.678 grams, so the microstructural counterpart of some folk-psychological generalization makes certain predictions regarding non-human creatures that cannot be justly anticipated solely from the microstructural counterpart's instances involving human beings, since the physical realizations of the same intentional states in non-human creatures cannot be extrapolated from these instances.

Second: Consider two possible worlds that are identical in all their Humean facts (roughly speaking: in their singular, local, non-nomic facts); in particular, suppose that in each, there is nothing but a single lonely proton moving with a given uniform velocity. One of these worlds is what would have been, had there been nothing but such a lonely proton; the laws of nature would have been no different from what they actually are.<sup>28</sup> For example, Coulomb's law would still have been a natural law, so had there also been an electron in this world, its electrostatic interaction with the proton would have been just as it is in the actual world. We might specify that in the other lonely-proton world, some alternative to Coulomb's law is in force. The laws, then, fail to supervene on the Humean facts; these two possible worlds agree in all of their Humean facts but differ in their natural laws. Nomic non-supervenience is incompatible with many proposed accounts of natural law but would be explained by lawhood's connection to salience. If, for any law, there must be some perspective from which evidence can render it salient to us, then the laws in a given possible world cannot be fixed by the Humean facts there, since our justificatory practice (a normative matter: what claim is rendered salient by what evidence regarded from a certain outlook) is not fixed by the Humean facts in a given possible world (even if it is the actual world). Rather, any supposition concerning our justificatory practice is compatible with the Humean facts in a given possible world. Consequently, once we have stipulated from scratch a certain possible world's Humean facts (as I did in specifying the second lonely-proton world above), we are free to stipulate any of

the regularities holding of these facts as expressing laws in that world.

Third: In Sellars's argument, the empirical-layer counterpart of a theoretical law is salient from neither the empirical outlook nor the theoretical outlook, yet it is a logical consequence of natural laws. Likewise, the microphysical counterpart of a folk-psychological generalization follows from the laws of physics but is not salient from a microphysical outlook. Lawhood's link to salience entails, then, that some nomically necessary generalizations (i.e., some logical consequences of law-statements) do not express natural laws. This result,<sup>29</sup> though contrary to many accounts of natural law, derives independent plausibility from scientific practice. Scientists have regarded certain generalizations as nomically necessary but coincidental rather than lawful, and we can understand their motivations for doing so. For example, some nineteenth-century chemists held that whereas it is a law that the molecular weights of the noncyclic alkanes differ by multiples of 14 units, and it is a law that the atomic weight of nitrogen is 14 units, it is merely a (nominally necessary) coincidence that the molecular weights of the noncyclic alkanes differ by multiples of the atomic weight of nitrogen (see van Spronsen, 1969, pp. 73ff.). Just as the microphysical counterpart of a folk-psychological law involves instances too disparate (considering the diversity of an intentional state's microphysical realizations) for examined instances to suggest the unexamined instances, so the fact that each examined nitrogen atom is 14 mass units fails to suggest a claim regarding the noncyclic alkanes; grouping the nitrogen atoms together with the noncyclic alkanes appears just as gerrymandered, as arbitrary, as grouping together all the microphysical realizations of a given folk-psychological state appears from a physical standpoint.

Thus, there are several reasons for thinking that there may be a connection between a regularity's lawfulness and its salience. What conception of lawhood would ground such a connection? Let me sketch one possibility, leaving its detailed elaboration for another occasion.<sup>30</sup>

A basic presupposition of scientific research is that we do not need to examine everything in order to know everything. Rather, a few observations, restricted in space, time, and other respects,

sometimes suffice to render salient a hypothesis that is accurate to all unexamined cases in a remarkably wide range of conditions. When does this happen? As we saw in connection with Regnault's and van der Waals's equations, the same observations considered in different frameworks may suggest different hypotheses. And it may be that some generalization is not rendered salient in a given framework even by many of its instances, as (Sellars argues) is the case for the phenomenal counterpart of a material-object law considered from a lower-layer (autonomously phenomenal) standpoint. On the other hand, it may take very few instances to render a certain generalization salient in a given framework. We are best off employing an "inductive strategy" (recalling this term from section I) that results in our projecting a maximally informative generalization on the basis of a minimum of observations. (Of course, prior to any empirical investigation, we may well not know what strategy this is.) As Mill puts the point:

[T]here are weighty scientific reasons for giving to every science as much of the character of a Deductive Science as possible; for endeavoring to construct the science from the fewest and the simplest possible inductions, and to make these, by any combinations however complicated, suffice for proving even such truths, relating to complex cases, as could be proved, if we chose, by inductions from specific experience. Every branch of natural philosophy was originally experimental; each generalization rested on a special induction, and was derived from its own distinct set of observations and experiments . . . . [A]ll these sciences have become to some extent, and some of them in nearly the whole of their extent, sciences . . . whereby multitudes of truths, already known by induction from as many different sets of experiments, have come to be exhibited as deductions or corollaries from inductive propositions of a simpler and more universal character . . . . A science is experimental, in proportion as every new case, which presents any peculiar features, stands in need of a new set of observations and experiments – a fresh induction. It is deductive, in proportion as it can draw conclusions, respecting cases of a new kind, by processes which can bring those cases under old inductions . . . (1893, p. 164; Bk. II, ch. 4, sect. 5).

Let's try to cash this out a bit.

For instance, suppose we search for overly narrow generalizations – say, generalizations specifying the physical properties of copper objects in North America. Then although accurate generalizations (e.g., "All copper objects in North America have a melting point of 1,083 °C") quickly become salient, we miss out on the benefits of using observations of North American copper

objects to suggest generalizations accurate to other copper objects. We would not be misled by projecting the salient regularities onto these other cases. We do not need to make a “fresh induction” – to carry out a separate inductive strategy – for copper objects outside North America; we do not need “a distinct set of observations and experiments” – i.e., we do not need to observe copper objects outside North America – in order to arrive at a generalization that is accurate to these cases. Our justificatory practices (i.e., what we should judge, on the given framework, would make unexamined copper objects outside North American *like* the North American copper objects we have examined) would lead us to generalizations accurate to copper objects outside North America, even though our observations consist exclusively of North American copper objects.

The opposite disadvantage results from pursuing an inductive strategy that seeks overly broad generalizations – say, generalizations covering the physical properties of objects that are either pure copper or pure nickel. Observations of copper objects will mislead us regarding nickel objects (and vice versa); the regularities covering all copper-or-nickel object that are suggested by observations exclusively of copper objects typically do not hold of nickel objects. We need to make a “fresh induction” for nickel objects. In the same way, suppose we are seeking a microphysical generalization covering all instantiations of a given intentional state, and suppose that all of the cases we have observed of this intentional state are of its realization in human beings. Then these observed instances will mislead us regarding the microphysical realizations of this intentional state in other creatures. We need a fresh induction for each kind of creature, whereas we do not need such a fresh induction if we proceed from a folk-psychological perspective; observations of human beings alone will not mislead us if we proceed within that framework.

Considerations like these might make one set of inductive strategies the best set for us to pursue in order to learn about the actual world. Not every general truth would be generated by pursuing the best set of inductive strategies. For instance, in carrying out an inductive strategy on which the law of thermal expansion becomes salient, we would regard instances of this law as supporting

its projection onto unexamined cases, actual and counterfactual, and so as justifying the counterfactual “Had a copper object of nearly a cubic mile in volume been heated, then it would have expanded to exceed a cubic mile.” By adopting this counterfactual, we are precluded from pursuing an inductive strategy that results in the projection of the generalization “All copper objects are smaller than a cubic mile”. That is because this generalization’s projection would include believing that a counterfactual copper object that is nearly a cubic mile and heated would behave in the same way as examined copper objects – i.e., would be smaller than a cubic mile.

Thus, the best set of inductive strategies for us to pursue does not yield all of the general truths. If the natural laws are identified with the regularities yielded by the best set of inductive strategies for us to carry out, then lawhood would be essentially connected to salience in the manner demanded by the three analogous arguments I have examined. On this view, the laws in a given possible world are associated with the results of the best set of inductive strategies for us to pursue in order to learn about that world – the best for *us*, here in the actual world, to carry out. We cannot learn about the closest lonely-proton world (i.e., about what would have been, had there been nothing but a single lonely proton moving uniformly forever) by examining that world directly; it is not available for our direct examination. Our only epistemic access to that world is through the actual world – in other words, by using observations of the actual world to confirm various counterfactual conditionals. So the best set of inductive strategies for us to pursue, in order to learn about the closest lonely-proton world, is the best set of inductive strategies for learning about that possible world on the basis of evidence drawn from the actual world. Now if we believe that some generalization expresses a natural law of the actual world, then (according to the account I am sketching) we believe it to be generated by one of the inductive strategies we should use in order to learn about the actual world, and so we treat the generalization as having been inductively projected from our evidence. Therefore, we believe that this generalization would have obtained, had there been nothing but a lonely proton, since this is one of the hypothetical unexamined cases to which it has been projected. So we regard the best set of induc-

tive strategies for us to carry out, in order to learn about the actual world, as leading to the truth about the closest lonely-proton world. We must regard any other set of inductive strategies as either leading to inaccurate conclusions concerning the closest lonely-proton world, or as arriving at accurate conclusions less efficiently than does the best set of inductive strategies for learning about the actual world. So we believe the best set of inductive strategies for us to pursue, in order to learn about the closest lonely-proton world, to be the best set of inductive strategies for us to pursue, in order to learn about the actual world. For this reason, we take the laws in the closest lonely proton world to be the same as the laws in the actual world.<sup>31</sup> Notice, then, that on this view, the natural laws in a given possible world are essentially connected to *us*, and this connection is crucial to accounting for a basic feature of scientific reasoning.

A connection between a regularity's lawfulness and its salience from some outlook would make it possible for ethics and folk psychology to be disciplines independent of physics, discovering laws of nature (as well as explanations, counterfactuals, and other law-involving matters of fact) that are irreducible to the laws discovered by physics. The moral and folk-psychological laws would constitute connections in the world not captured by the underlying physics. This would make moral and folk-psychological properties just like the properties discovered by physics in contributing novel laws to the world.<sup>32</sup> Sellars's argument against phenomenalism and scientific anti-realism, then, may contain the key insight that undermines the layer-cake pictures of science, ethics, folk psychology, and other disciplines.

#### ACKNOWLEDGMENT

My thanks to Jim Bogen, Paul Hurley, John McDowell, Ruth Millikan, and Seana Shiffrin for comments on earlier drafts. I am also deeply grateful to Robert Brandom, who introduced me to Sellars's work.



## NOTES

<sup>1</sup> Sometimes, the “layer cake” is depicted as having a third, even lower layer, consisting of observation reports.

<sup>2</sup> There are other important arguments having the very same structure. See, for example, Kitcher (1984, esp. pp. 349ff., 370ff.).

<sup>3</sup> Sellars (1963c) rejects the view that our observation reports consist of such claims as “It looks like there is a red object before me now.” But when Sellars argues against the layer-cake picture, he concedes to phenomenalism its conception of our observation reports.

<sup>4</sup> Here is how Van der Waals’s equation can be derived by amending the quantities in the ideal gas law in light of molecular considerations. The “ideal” volume per mole is the volume that a mole of gas would occupy if a molecule itself had no size; the “ideal” volume per mole equals the measured volume per mole minus some constant reflecting a molecule’s size. Van der Waals thus replaced  $V$  by  $(V-b)$ . The “ideal” pressure is what the pressure would be if there were no intermolecular attractive forces; it equals the measured pressure plus some quantity reflecting the extent to which an individual gas molecule hitting the container wall is slowed by being attracted to gas molecules behind it, and so hitting the wall with diminished force. The back-attraction felt by a single molecule is proportional to the density of the gas behind it, and so to  $1/V$ . The back-attraction’s effect on the pressure is its effect on a single molecule multiplied by the frequency of molecular collisions with the container wall, which is again proportional to the density of the gas, and so to  $1/V$ . So Van der Waals replaced the “ideal” pressure by the measured pressure plus some constant (reflecting the strength of intermolecular attraction) multiplied by  $1/V^2$ . See my (2000, pp. 212–20).

<sup>5</sup> I cash out this notion of “inductive confirmation” more carefully in my (1996, 1999 and especially 2000).

<sup>6</sup> This is an instance of the claim, familiar since Goodman and Reichenbach, that we can project a hypothesis only if we believe that it may state a law. Van Fraassen (1989) rejects this view; I refine and defend it in my (1996, 1999 especially 2000).

<sup>7</sup> I have omitted many details of this argument; I treat it more carefully in Lange (1996).

<sup>8</sup> McDowell (1981, p. 144) asserts that the moral supervenes on the naturalistic (or on what is typically called “naturalistic,” since for McDowell, the genuinely naturalistic properties include the moral properties). But “supervenes” in what sense? McDowell does not say. Blackburn holds that the moral supervenes on the naturalistic in the sense that two naturalistically indistinguishable entities in a given possible world must be morally indistinguishable. This sense of supervenience allows two entities, in distinct possible worlds, to be naturalistically indistinguishable but morally distinct. Blackburn holds that to call a clearly good thing “bad” is to make a moral error rather than to suffer from a conceptual confusion. Therefore, different distributions of moral properties are *logically* compatible with the same distribution of naturalistic properties. This view requires a sharp distinction between understanding moral predicates and

knowing how to apply them correctly, and correspondingly between logical necessity and a weaker modality, “moral necessity,” holding between (say) an action’s naturalistic properties and the moral evaluation necessitated by them. This distinction seems unwarranted in view of the demise of the analytic/synthetic distinction. Accordingly, by the “supervenience” of the moral on the natural, I mean something more like this: any two entities, whether in the same possible world or in any two possible worlds, that are naturalistically indistinguishable are morally indistinguishable. (This departure from Blackburn shall not affect my discussion of his criticism of McDowell’s arguments against projectivism, though Blackburn’s positive arguments for projectivism depend on his embracing the weaker sense of “supervenience” rather than the “cross-world” sense.) But what sense of “possible” do I mean when referring to “any two possible worlds” in the construal of “supervenience”? Let me put this question aside for the moment; I return to it in section III, when I discuss the sense in which a supervening predicate is “non-accidentally coextensive” (i.e., coextensive in all possible worlds) with some subvening predicate.

<sup>9</sup> That this kind is not finitely expressible is irrelevant. Such a non-accidentally coextensive natural kind is guaranteed by the stronger, “cross-world” notion of supervenience, not by the weaker, “intra-world” conception; see note 8.

<sup>10</sup> There is much in Blackburn’s remark with which I agree, as I discuss shortly.

<sup>11</sup> Wiggins (1987, pp. 199f.) agrees with McDowell on these points and is clear on this conjunction.

<sup>12</sup> Wiggins puts the point well: A justification of our judging cruel the hoodlums’ setting fire to the cat

will have to place our particular reaction in the face of what the hoodlums did in some relation to our responses to a whole heterogeneous collection of events whose only common property is that they are cruel in some degree (or so one is tempted to put it. What else after all do the following have in common? – A man snubs a child; the courtiers steal Rigoletto’s daughter; the hoodlums set fire to the cat ...) ... (1987, p. 157).

This reveals one source of the apparent cogency of Blackburn’s remark about the comic. We could in principle justify judging a given thing to be comic in the way Wiggins has us justify judging something to be cruel. But typically we don’t offer any argument to justify judging something to be comic. Since McDowell’s argument (as I see it) concerns reasoning – i.e., pertains to the context of justification – Blackburn’s analogy to “comic” is misleading. Likewise, I steer away from McDowell’s (1985) comparison of moral properties to secondary qualities; while the sense in which secondary qualities are objective may be a useful model for moral properties, the comparison is potentially misleading since typically, we don’t give an argument to justify “This is red” any more than we do for “This is funny”. These comparisons, then, de-emphasize salience’s normative role, making salience seem like a mere psychological matter.

<sup>13</sup> For more on Sellars' comparison of induction to the Socratic elenchus, see Brandom (1988).

<sup>14</sup> It is sometimes thought that McDowell is a particularist – that he denies the possibility of any such moral generalizations. This interpretation is encouraged by his talk of “uncodifiability,” for instance, but I don’t believe it is strictly accurate. Rather, as I am about to explain, he denies that there are (many?) moral generalizations having “independent intelligibility,” i.e., lacking *ceteris-paribus* clauses that can be arrived at by reasoning autonomously from outside the moral outlook. At least, this is the view with which I shall be concerned; anything more particularist threatens supervenience.

<sup>15</sup> One of the main points made by Kuhn (1970) is that learning a scientific theory (especially a “paradigm”) involves not only learning certain putative laws, but also learning certain skills of recognition and categorization (as well as certain values and expectations).

<sup>16</sup> For a fuller discussion of this interpretation of *ceteris-paribus* clauses in statements of natural law, see Lange (1993, 2000).

<sup>17</sup> In the present paper, I do not distinguish concepts from properties or predicates. Also, there are some delicate issues about what qualifies as a “physical predicate”, and these affect the plausibility of my claim; see Heil (1992, pp. 71ff.).

<sup>18</sup> This is part of Blackburn’s point in the passage I quoted earlier. Note that even if the best explanation of someone’s moral evaluations were that she is tracking certain moral properties really possessed by those things, this would not suggest to van Fraassen that those things really possess those moral properties. According to van Fraassen, explanatory power over and above empirical adequacy is not indicative of truth (over and above empirical adequacy).

<sup>19</sup> Wiggins (1987, pp. 158ff.) might be taken as making the argument I am rejecting here. For more, see Lange (1996).

<sup>20</sup> McDowell (1985, p. 129) explains that part of his disagreement with Blackburn is over whether the “justifications” of particular evaluations should be understood as genuine justifications or mere “quasi-justifications.” McDowell cannot presume that they are justifications without begging the question, even though if they are not justifications, then it is “problematic whether evaluative language is close enough to the usual paradigms of concept application to count as expressive of judgments at all (as opposed to a kind of sounding off)” (McDowell 1981, p. 158).

<sup>21</sup> Note, again, that this point concerns the context of justification.

<sup>22</sup> That intentional states function in science much like the unobservable entities posited by physical theories is the underlying motif of Sellars’s “Myth of Jones” in (1963c).

<sup>23</sup> Of course, “injury” (for instance) is not an intentional state. But a law of folk psychology can mention other matters besides intentional states; its doing so need not threaten to turn it into a law of microphysics. (Schiffer (1991, p. 5 n. 2) apparently agrees.) “Injury” is certainly not mentioned in a microphysical law.

<sup>24</sup> The comparison of folk-psychology to geometrical optics is especially apt in light of the fact that according to many, both branches of knowledge are nearly complete.

<sup>25</sup> This talk of “folk-psychological laws” does not diminish Davidson’s point that “the constitutive ideal of rationality” shapes our thought about propositional attitudes (1990, p. 223). The possibility of such folk-psychological laws is compatible with the view that what it is for someone to believe that *p* (e.g., her undertaking various commitments and responsibilities in virtue of the belief’s content) cannot be reduced to anything non-intentional or non-normative. Folk-psychological laws do not specify what it is for someone to believe that *p*, but what is nomically associated with so believing; folk-psychological laws characterize what does happen, not what ought to happen. An explanation of an intentional state or action that explains why it is rational is distinct from an explanation that explains why it happens, though both explanations are intentional. Moreover, it may be that folk-psychological generalizations are salient only from a perspective that recognizes certain patterns of behavior as required by ideal rationality, so that from this perspective, the hypotheses that become salient concern the circumstances in which agents depart from these patterns.

<sup>26</sup> For further discussion of “autonomy” in roughly this sense, see Hacking (1990, pp. 180ff.), Kitcher (1984, pp. 370–372), and Putnam (1975, p. 296). This notion of “autonomy” seems intended to capture two related ideas: (i) that these domains are self-lawed, i.e., governed by their own distinct laws, and (ii) that these laws would still have held, even if certain details from other levels had been different. These two ideas are related because laws are distinguished from accidental generalizations by their capacity to support counterfactuals. See Lange (2000).

<sup>27</sup> Certain “necessities of psychology” may be “accidents of physics”, such as (perhaps) “Practice improves performance, *ceteris paribus*.” This may indeed be a consequence of the neurological architecture we happen to possess.

To my line in sections II–III, one might object that folk-psychological (and moral) “laws” are not natural laws at all, but analytic truths. This is not to construe them as trivial in the same way as “All *F*’s are *G* except when they aren’t,” but as expressing what we mean by “believing that *P*” (or “being courageous”). I cannot address this issue fully here, but see note 25. See also Churchland (1970, p. 225), who emphasizes the demise of the traditional notion of analyticity. Furthermore, since the *ceteris-paribus* clause does not trivialize a folk-psychological generalization, it is not evident to me that the generalization is analytic even as analyticity is traditionally understood. Isn’t it a contingent matter, not obtaining in all logically possible worlds, that agents (ourselves and others) get confused, misremember (or even have memory), have parts of their bodies on whose health their minds depend, etc.?

<sup>28</sup> Scientific practice bears out this intuition. Cosmologists, for example, run computer simulations of what the universe’s history would have been under various different initial conditions. The computer program governs the simulation in accordance with what the cosmologists believe to be the natural laws; when cosmologists vary the initial conditions they feed into the program, they leave

the program itself unchanged. To test their programs, cosmologists sometimes set the initial conditions so that there is nothing forever but a single lonely proton, and they hold the program (the laws) fixed under this supposition. For more, see Lange (2000, Ch. 2).

<sup>29</sup> That lawhood fails to be closed under logical consequence is explicitly embraced by Fodor (1974, pp. 109ff.), but he is among the very few who do so (compare, for example, Davidson 1990, p. 218). It is unclear to me whether Sellars recognizes this as a consequence of his argument. The closest he comes to doing so is (1977, p. 313), but other passages (e.g., 1977, p. 311) suggest that he regards the empirical counterparts as lawful, though not ascertainable by working autonomously in the empirical framework.

<sup>30</sup> I have sketched this possibility elsewhere; see Lange (1995). I develop it systematically in Lange (2000).

<sup>31</sup> Of course, by inductively confirming that all copper objects have a melting point of 1,083 °C, we learn nothing about the copper objects in the closest lonely-proton world, since by stipulation, there are none. But we do learn *something* about the closest lonely-proton world: we learn what copper objects *would* have been like there, had there been any. In other words, when we inductively confirm the above generalization about copper objects, we confirm various nested counterfactuals, such as “Had there been nothing but a lonely proton, then had there been a copper object, it would have had a melting point of 1,083 °C.”

<sup>32</sup> For example, the electric field’s properties at a given location figure in special laws.

## REFERENCES

- Blackburn, S. (1981): ‘Reply: Rule-Following and Moral Realism’, in S. Holtzman and C. Leich (eds.), *Wittgenstein: To Follow A Rule*, London: Routledge, pp. 163–187.
- Blackburn, S. (1993): ‘Realism, Quasi, or Queasy?’ in J. Haldane and C. Wright (eds.), *Reality, Representation, and Projection*, New York: Oxford University Press, pp. 365–384.
- Brandom, R. (1988): ‘Inference, Expression and Induction: Some Sellarsian Themes’, *Philosophical Studies* 54, 257–285.
- Churchland, P. (1970): ‘The Logical Character of Action Explanations’, *Philosophical Review* 79, 214–236.
- Davidson, D. (1990): *Essays on Actions and Events*, Oxford: Clarendon.
- Fodor, J. (1974): ‘Special Sciences’, *Synthese* 28, pp. 77–115.
- Hacking, I. (1990): *The Taming of Chance*, Cambridge: Cambridge University Press.
- Heil, J. (1992): *The Nature of True Minds*, Cambridge: Cambridge University Press.
- Hempel, C. (1965): ‘The Theoretician’s Dilemma’, in *Aspects of Scientific Explanation*, New York: Basic Books, pp. 173–228.

- Kitcher, P. (1984): '1953 And All That: A Tale of Two Sciences', *Philosophical Review* 93, 335–373.
- Kuhn, T. (1970): *The Structure of Scientific Revolutions*, 2nd ed., Chicago: University of Chicago Press.
- Lange, M. (1993): 'Natural Laws and the Problem of Provisos', *Erkenntnis* 38, 233–248.
- Lange, M. (1995): 'Are There Natural Laws concerning Particular Biological Species?', *Journal of Philosophy* 92, 430–451.
- Lange, M. (1996): 'Laws of Nature, Cosmic Coincidences, and Scientific Realism', *Australasian Journal of Philosophy* 74, 614–638.
- Lange, M. (1999): 'Why Are the Laws of Nature so Important to Science?', *Philosophy and Phenomenological Research* 59, 625–52.
- Lange, M. (2000): *Natural Laws in Scientific Practice*, New York: Oxford University Press.
- Loeb, L. (1934): *The Kinetic Theory of Gases*, New York: McGraw-Hill.
- McDowell, J. (1978): 'Are Moral Requirements Hypothetical Imperatives?', *Proceedings of the Aristotelian Society* supplementary volume 52, pp. 13–29.
- McDowell, J. (1979): 'Virtue and Reason', *The Monist* 62, 331–350.
- McDowell, J. (1981): 'Non-Cognitivism and Rule-Following', in S. Holtzman and C. Leich (eds.), *Wittgenstein: To Follow A Rule*, London: Routledge, pp. 141–162.
- McDowell, J. (1985): 'Values and Secondary Qualities', in T. Honderich (ed.), *Value and Objectivity*, Oxford: Oxford University Press, pp. 110–129.
- McDowell, J. (1987): 'Projection and Truth in Ethics', *The Lindley Lecture*, University of Kansas.
- Mill, J.S. (1893): *A System of Logic*, 8th ed., New York: Harper and Bros.
- Millikan, R.G. (1993): 'Explanation in Biopsychology', in J. Heil and A. Mele (eds.), *Mental Causation*, Oxford: Clarendon, pp. 211–232.
- Putnam, H. (1975): *Mind, Language, and Reality*, New York: Cambridge University Press.
- Schiffer, S. (1991): 'Ceteris Paribus Laws', *Mind* 100, 1–17.
- Sellars, W. (1963a): 'Phenomenalism', in *Science, Perception and Reality*, London: Routledge, pp. 60–95.
- Sellars, W. (1963b): 'The Language of Theories', in *Ibid*, pp. 106–126.
- Sellars, W. (1963c): 'Empiricism and the Philosophy of Mind', in *Ibid*, pp. 127–196.
- Sellars, W. (1965): 'Scientific Realism or Irenic Instrumentalism', in R. Cohen and M. Wartofsky (eds.), *Boston Studies in the Philosophy of Science*, v.2, New York: Humanities Press, pp. 171–204.
- Sellars, W. (1977): 'Is Scientific Realism Tenable?' in *PSA 1976 v.2*, Lansing, Michigan: Philosophy of Science Association, pp. 307–334.
- Teller, P. (1984): 'Comments on Kim's Paper', *Southern Journal of Philosophy* 22 supp. vol., pp. 57–61.
- Wiggins, D. (1987): *Needs, Values, Truth*, Oxford: Blackwell.

- van Fraassen, B. (1977): 'On the Radical Incompleteness of the Manifest Image', in *PSA 1976* v.2, Lansing, Michigan: Philosophy of Science Association, pp. 335–343.
- van Fraassen, B. (1980): *The Scientific Image*, Oxford: Clarendon.
- van Fraassen, B. (1989): *Laws and Symmetry*, Oxford: Clarendon.
- van Spronsen J.W. (1969): *The Periodic System of Chemical Elements*, Amsterdam: Elsevier.

*Department of Philosophy*  
*University of Washington*  
*345 Savery Hall*  
*Box 353350*  
*Seattle, WA 98195–3350*  
*USA*

