"Saving the Phenomena" Today

Paul Teller†

Bogen and Woodward argued the indirect connection between data and theory in terms of their conception of "phenomena." I outline and elaborate on their presentation. To illuminate the connection with contemporary thinking in terms of models, I distinguish between phenomena tokens, representations of which can be identified with data models, and phenomena types that can be identified with relatively low-lying models or aspects of models in the model hierarchy. Throughout I stress the role of idealization in these considerations.

1. Introduction. The expression "saving the phenomena" goes back to ancient Greek astronomy and has been variously used ever since. Apparently twentieth-century usage was much influenced by Duhem's discussion (1969), in which he employed the phrase by way of attributing an instrumentalist attitude toward the ancient Greek project of characterizing the positions of the "wandering stars." In a much-read and cited paper of 1988, James Bogen and Jim Woodward (henceforth B&W; uncited page numbers are all to this paper) adopted the phrase for their own, very different use. In the following pages I will summarize the view that they presented, offer some small clarifications, and then offer my own suggestion for recharacterizing what I take to be their ideas in a more contemporary idiom.

2. Data, Phenomena, and Theory. The 1960s and 1970s critique of positivism greatly liberalized the positivists' view of what might count as an observation, but when B&W wrote in 1988, the prediction of individual observations seemed, as it largely does today, a central method of theory evaluation and theory use.

B&W reject this view: "The underlying [positivist] idea that scientific theories are primarily designed to predict and explain claims about what

†To contact the author, please write to: 1 Shields Avenue, University of California, Davis, CA 95616; e-mail: prteller@ucdavis.edu.

Philosophy of Science, 77 (December 2010) pp. 815–826. 0031-8248/2010/7705-0002$10.00
Copyright 2010 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/656537 Published online by Cambridge University Press

we observe remains enormously influential, even among the sharpest critics of positivism. . . . Our aim in this paper is to show that this aspect of the traditional account is fundamentally incorrect" (204–5). B&W set out an alternative way of thinking about the relation between theory and individual observations—what they refer to as "data"—by appealing to what they called "phenomena": "Our general thesis is that we need to distinguish what theories explain (phenomena or facts about phenomena) from what is uncontroversially observable data" (314).

To illustrate, consider the prediction of general relativity of the deflection of starlight passing near the sun. "Raw data" collected for confirmation, the locations of individual spots on photographic plates, were hardly themselves predicted by general relativity! The theory predicted the phenomenon of bent rays of light, not the individual spots on photographic plates. Other examples of phenomena that B&W examine in some detail include the melting point of lead, weak neutral currents, the rate of neutrino emission from the sun, and various functional characteristic of our brains, such as behavior preservation and various features of sensory processing.

Today many would describe some of these cases using the idea of a data model. Later on I will examine the relation between B&W's approach; the idea of data models; and, more broadly, contemporary ways of thinking about the role of models in the scientific enterprise.

B&W provide some general characteristics of what they have in mind when they talk about "phenomena." A phenomenon can occur in many situations with stable, repeatable characteristics (317). A phenomenon results from the interaction of some manageably small number of causal factors (317). It is phenomena, not data, that are explained by theories (305–6, 322, 326, 335). Phenomena serve as evidence for theories (306) and are themselves generally not observable (306, 245, 347).[1] Finally, B&W are explicitly neutral about the ontology of phenomena. Their phenomena belong to the "natural order" but can be objects, features, processes, states, and so on (321–22).

By "data" B&W intend individual events, or their representations, that one might think of as of the written-down-in-the-lab-book sort. Individual events that serve as the material for statistical analyses constitute a special case. B&W say that data have complex causes and are idiosyncratic to particular experimental situations (317). Data are generally observable (312). Data provide evidence, not for theory but for phenomena. And to

1. "We use 'observe' to mean 'perceive' or 'detect' by means of processes which can be usefully viewed as extensions of perception" (305). The sentence I have quoted occurs in the antecedent of a conditional, but context clearly indicates that B&W mean to endorse it (see 342–52).

https://doi.org/10.1086/656537 Published online by Cambridge University Press

serve this evidential role, in practice data must be accessible in sufficiently large quantities to provide evidence (320) and must be tractable for data reduction and statistical analysis (321).

In addition, B&W specify that "data must also be such that it is relatively easy to identify, classify, measure, aggregate, and analyze in ways that are reliable and reproducible by others" (320). This does not mean that the data, which are "idiosyncratic" themselves, need to be reproducible, one by one. Rather, the requirement of reproducibility is really a requirement on phenomena: the data must be such that others can find or produce the same kind of data, that is, data providing similar evidence for the same phenomenon.[2]

Let me mention some qualifications and clarifications. Rather than thinking of their characterizations of phenomena and data as constituting essentializing necessary and sufficient conditions, we should think of the characterizations as "ideal types" to which real world instances approximate more or less closely: "there really is no principled [that is, no absolutely sharp] distinction to be drawn between data and phenomena" (315). Throughout, they use the qualifications "typically," "for the most part," and "in most cases," though these qualifications are sometimes dropped.

As mentioned, B&W suggest that phenomena are individuated by their causes: "the occurrence of these instances [of a phenomenon] is (or is plausibly thought to be) the result of the interaction of some manageably small number of causal factors, instances of which can themselves occur in a variety of different kinds of situations (317).[3] Thinking of phenomena as individuated by their causes makes it easy to understand what it is that we learn when we learn more about a phenomenon of which our understanding was initially superficial.

A simple illustration will show how natural it is to take phenomena to be individuated by their causes. A whistling train rushes by you, and as it passes, the pitch of the whistle suddenly drops: an instance of the Doppler effect. Think now, instead, of a train sitting by with a whistle tone that suddenly drops, exactly as much as for the moving train, but in the second case the change is due to the engineer manually adjusting

2. The intended reading, Jim Woodward tells me.

3. There is an issue about how "causes" is to be understood. The full causal background of most events will be enormously complex, precisely what B&W associated with data. B&W appear to take for granted a notion (or notions) of cause exemplified by a match being struck causing it to light, smoking causing cancer, and neutron absorption causing fission, neither offering nor referring to any analysis. The worry is that causal connections in such a sense might themselves be taken to be phenomena, as in these three examples. Woodward's later analysis of cause (2003, 2007) might fill the gap here.

https://doi.org/10.1086/656537 Published online by Cambridge University Press

the whistle pitch. The effect sounds the same, but different cause, different phenomenon.

Individuating by causes gives natural further detail in the relation among phenomena. Diabetes is a phenomenon caused by insufficient insulin uptake. One cause, one phenomenon. But this cause itself has two different causes: insufficient production of insulin, type 1 diabetes, or resistance to its effects, type 2 diabetes. The one phenomenon, diabetes, can be analyzed into two more specific subphenomena.

3. Some Friendly Amendments. Lots of phenomena are plainly observable in the common sense of being perceptible: tea kettles humming, the Doppler effect, Poisson spots, eclipses, optical illusions generally, iron filings lined up by a magnet, radiometers, potato batteries, the blue sky, rising bread, the double slit experiment, and so forth. But B&W did not claim that all phenomena are unobservable, and despite the obvious length of the list that we could extend, I do not think their attitude toward phenomena as typically unobservable is terribly wide of the mark. Note that most of the items on my list are from our everyday experience or from older chapters in science. As the sciences mature and become more complex, they increasingly focus on unobservable phenomena. The way in which a science increasingly works in terms of unobservable phenomena might serve as an interesting parameter in the development and characterization of a science.

Need data be copious? Not always. If Galileo really did perform the Power of Pisa experiment, he would not have had to repeat it many times. One can satisfy oneself that a potato battery produces a current with just one or, at most, a very few tests. In contrast to the high-energy experiments that B&W cite, some discovered phenomena in high-energy physics were established on the basis of just a very few "golden events." But B&W qualified their claim: "Matters must be arranged so that data is produced sufficiently frequently and in sufficiently large amounts that human beings can detect enough of it in reasonably short periods of time to support conclusions about the existence of phenomena" (320). It is consistent with this that in special cases just one datum or a small number of data will suffice.

Still, B&W may seem to overstate the case when they write that "it is hard to imagine an instrument which could register any phenomenon of interest in isolation from the multitude of background factors which operate in such a way as to make data-analysis and interpretation necessary" (352). Does use of standard instruments, such as voltmeters, provide counterexamples? Only superficially. We take the operation of well-designed instruments for granted, but getting that good design right was usually a difficult and arduous task (Gooday 1995). Tacitly, the use of a simple

https://doi.org/10.1086/656537 Published online by Cambridge University Press

voltmeter rests on just the kind of arduous work that B&W have in mind, where this work has been embodied in the instrument and its design.

The last quote suggests the metaphor of thinking of data analysis as sifting out "noise" so we can hear the "signal." Messy bits of data arise from a tangle of causal factors, so we need data analysis to filter the signal, constituted by the causal component contributed by the phenomenon of interest, from the noise, constituted by other causal features that obscure the effects of the phenomenon of interest. Other essays in this symposium will discuss ramifications of this metaphor, so it may be useful for me to sketch a brief introduction.

An individual data point is produced by a tangle of causal considerations. One of these, or a small coherent collection, will count as the signal, the rest as obscuring noise. But which is which? The suggestion has been that the signal comes from the phenomenon that one is trying to identify. If so, since what phenomenon one is looking for depends on interests, what counts as signal and what counts as noise can shift with change of interests.[4]

To illustrate with a stylized example: lightning produces static in AM radio broadcast signals. Imagine that by carefully analyzing such static we could extract useful information about lightning strikes—how far away, how strong. But now, when considering an AM radio broadcast in a thunderstorm, what constitutes the signal, what the noise? In ordinary situations the causal component generated by the talk that is being broadcast counts as the signal, while the causal component that comes from lighting strikes counts as static—as noise. But if our interest is in the lightning strikes, it is just the other way around. The features of the "static" caused by the lighting will be our "signal," while the component coming from the broadcasted program will be just interfering "noise."

What counts as problematic levels of noise is also relative to interests. One of B&W's examples of a phenomenon is the (value of) the melting point of lead—the fact that there is a stable temperature at which lead melts. But, say B&W, "one does not determine the melting point of lead by observing the result of a single thermometer reading" (308). Thermometer readings will be affected by causal factors other than the temperature one is trying to measure, so the individual thermometer readings—the data—must be subjected to assumptions about these obscuring

4. My perception of this idea was inspired by McAllister (2007). But, contrary to his current view (see his contribution to this symposium), I must emphasize that not any pattern that we project onto data will correspond to a phenomenon in nature. The patterns that count will only be those with a stable causal source that admit of repeated application. See further both Brading's (2010, sec. 3.2) and Woodward's contributions to this symposium.

https://doi.org/10.1086/656537 Published online by Cambridge University Press

causal features and then analyzed statistically to get an estimate and expected margin of error for the melting point.

Is such data analysis really required in this case? It depends! It depends on the accuracy that one wants. If not much accuracy is needed, one measurement with a well-calibrated thermometer will suffice. The operation of numerous other causes of variation or "error" are too small to matter. But if one requires greater accuracy, these small, perturbing causal factors must be taken into account and do count as problematic levels of noise.

The example of the melting point of lead brings up another consideration: there is no such thing as *the* melting point of lead. Environmental factors will affect the details of the ways in which lead turns from solid to liquid, and even for fixed environmental facts there will be no perfectly precise temperature at which melting occurs—indeed, the whole conception of temperature does not apply with absolutely complete precision.

So what, then, is the "phenomenon" in question in this example? It was described as "the melting point of lead," but *the* melting point of lead, the precise temperature at which melting occurs, is an idealization. So what will count as a phenomenon appears, at least in this example, to be tangled up with the ubiquitous use of idealizations in science. More on this below.

4. Token Phenomena, Type Phenomena, Idealization, and Regularity. It will be useful to distinguish between token and type phenomena. Consider the general fact that salt, when stirred into water, will dissolve. Contrast that with a specific event of some concrete quantity of salt being stirred into a specific glass of water. The later is a token of the former phenomenon type, an instance of a general fact.

Concrete phenomenon tokens present a potential puzzle. Any specific token event will be indefinitely complex in the details of its occurrence—in the foregoing example, the exact temperature and quantity of water in the glass, the height of the glass. So when we focus on the complex of details of a determinate phenomenon token, the phenomenon token looks to be every bit as idiosyncratic as data are supposed to be. Is there, after all, really a difference in this regard between data, on the one hand, and phenomena tokens, on the other?

To solve this puzzle remember that an individual event counts as a phenomenon token because it falls under a phenomenon type, and these types isolate some limited range of features in which we may be interested. In characterizing a concrete event, state, or process as a phenomenon, we are taking it as an instance of a phenomenon type, which type carries with it the implication of the narrow range of characteristics or processes that are in question.

https://doi.org/10.1086/656537 Published online by Cambridge University Press

B&W characterized phenomena as individuated by "the interaction of some manageably small number of causal factors" (317). I would think, however, that B&W would not want to limit what holds a phenomenon together to be a fixed set of *causal* influences. Consider the phenomenon that in reflection, the angle of incidence is equal to the angle of reflection. The mechanism by which this occurs varies widely from case to case—it holds as much for a bouncing ball as a beam of light. In the case of reflection it is a common geometrical fact that unifies the diverse processes into one kind of phenomenon.[5] What counts is that there be a basis, causal or otherwise, known or unknown, for regularity in the kind of occurrence that is in question. It will be the regularity and its basis pertaining to a phenomenon type in virtue of which a complex real world occurrence counts as an instance of a phenomenon.

We think of regularities as straightforward general facts about an independently existing world. But there are complications that surface with what Woodward (2007; see also 2003) calls "coarse-graining."[6] Consider the qualitative phenomenon constituted by the fact that a window struck by a flying rock will shatter. Woodward describes this phenomenon in terms of a pair of two-valued variables: Y, taking the value 1 (window struck by rock) and 0 (window not struck), and Z, taking the value 1 (window shatters) and 0 (window does not shatter). Variables Y and Z do not admit of any intermediate values (hit by an object that is a borderline case of a rock or hit at a grazing angle; intermediate between a clear case of shattering and not shattering), and values of X and Y will be variously realized by extremely different microstructures. Altogether, the generalization, if $Y = 1$, then $Z = 1$, is a shortcut idealization that is not exactly accurate.

What, then, is the status of the phenomenon that when a window is hit by a flying rock, it shatters? B&W say that "it should be clear that we think of particular phenomena as in the world, as belonging to the natural order itself and not just to the way we talk about or conceptualize that order" (321). There is no question, *tokens* of the phenomenon of rocks breaking windows are events in the natural order. But what of the phenomenon *type*? As an idealized type, it involves shortcuts and simplifications, so we have somehow gotten into the act along with nature.

To further illustrate with a case that we more immediately think of in terms of idealization, consider the phenomenon described by Hooke's law, $F = -kX$, where X is the extension of a spring, F is the restoring

5. Reversal of the dynamic vector component (momentum, propagation vector) perpendicular to the reflection surface.

6. For a quite different exposition of ideas very similar to what follows, see my article (Teller 2004).

https://doi.org/10.1086/656537 Published online by Cambridge University Press

force, and k is the spring constant of the spring in question. The variables F and X are macroscopic, conceived of as in classical physics. But just how values of these variables are realized by the underlying microstructure will be a complicated affair, differing from case to case. Moreover, as a quantitative idealization, Hooke's law applies to nothing exactly. To which materials does Hooke's law apply, even inexactly? Well, to springs. But there is no independent, general characterization of what will count as a spring: a spring is whatever Hooke's law (approximately) applies to.[7] It is an objective fact about nature that there is a diverse range of materials to which Hooke's law applies, in a range and to a degree of accuracy that we find useful. In fact, a great number of materials are described by Hooke's law but only in a range so small that it is of no interest to us. Altogether the phenomenon (type) characterized as Hooke's law is a complicated affair to which both we and nature contribute.

The same comments go for Boyle's law, the laws of reflection and refraction, and Euler's equations for the fluid behavior of liquids. The laws of fundamental theories such as quantum field theory and general relativity are not exempt. All exactly stated laws/regularities/phenomena known today are idealizations.

Do all phenomena (types) really involve idealization? What about the phenomenon of hydrogen and oxygen combining to form water? While this, and a great many other examples, involve no explicit idealization, there are no completely determinate things and their behaviors picked out by such descriptions that would constitute a precisely deliminable regularity to count as a phenomenon. If we try to characterize the components of this phenomenon in terms of atoms and their behavior, we are back to idealizations; we can steer clear of the idealizations only by withholding precision (just what counts as water? How much impurity? At what pressures and temperatures?). The use of imprecision—leaving what is in question in some ways open-ended—means that we again get into the act in determining just what counts as a phenomenon.[8]

I can do no more than call attention to these complications in the idea of phenomena. Sorting them out is an instance of very general puzzles about how we use vague truths and false idealizations to give us real knowledge of the world. But having located the idea of phenomena in

7. Hooke's law is the linear approximation to the Taylor series expansion of the function describing the material's distortion. We call something a spring when this approximation is good enough for our interests.

8. In fact, I believe that false, precisely stated idealizations and what I call their semantic alter egos, true but imprecise characterizations, are really two different representational tools that accomplish the same representational work. See my article (Teller 2008) and publications to come.

https://doi.org/10.1086/656537 Published online by Cambridge University Press

the circle of ideas concerning idealizations, let us look at how B&W's conception from 1988 fits into contemporary thinking about science as an activity of building and using idealized models.

5. Modeling and the Rediscovery of B&W's Phenomena/Data Contrast. During the past 2 decades the philosophy of science has developed an acute appreciation of the role of models in the practice of science and has looked in many ways at the roles that models play. The various sciences use different kinds of models in different roles. While there is no one fixed structure that fits all the sciences, at a sufficiently abstract level of description there is a general pattern that fits a great deal of scientific work. Different presentations differ in many details—I will present the idea by adapting the exposition in a recent paper by Ron Giere (2010) that I feel is fairly representative (see fig. 1).

A theory will deploy models that differ greatly in their level of generality. Taking physics as an example, and starting at the highest level of generality, Giere recommends that we think of Newton's laws not as general truths but as "model-building principles," stock building guidelines for building models. These principles are exemplified by "principled models" that do no more than instantiate the model-building principles. For example, a principled model based on Newton's three laws describes massive objects moving with unspecified velocities and accelerations, subject to unspecified forces, except as required by the satisfaction of the three laws.

Representational models result by characterizing principled modes in some more specific respects. For example, if we also characterize the force function as proportional to a displacement, we get "the harmonic oscillator," something that is still extremely abstract but intended to represent a more specific kind of system, such as a spring or pendulum. Such abstract representational models are made more and more specific. For example,

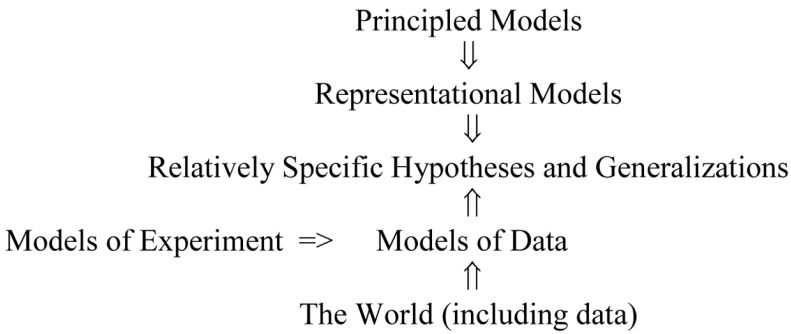


Figure 1.

https://doi.org/10.1086/656537 Published online by Cambridge University Press

we can elaborate the model by taking into account additional forces, such as the force of friction. We can become more specific about the spring or the pendulum that might be under consideration, for example, by specifying the spring constant or the length of the pendulum. This process of elaboration continues until finally we develop models intended to represent very specific, concrete, real systems.

So the top half of the hierarchy in the diagram will be elaborated with many layers, getting more and more specific as one goes down. In all of this, of course, one uses many tools of idealization and approximation, and Giere also notes that the more specific models in the hierarchy are often cobbled together from general principles from what we think of as very different theories.

Turning to the bottom half of the hierarchy, we also have a process of model building and elaboration. It is unusual for written-down-in-the-lab-book data to be compared directly with a specific hypothesis or generalization. Instead, we deploy general ideas and specific assumptions, collectively called "models of experiments," to interpret the data. For example, one may use well-developed statistical techniques along with further assumptions, such as that the data are normally distributed. We often use general theoretical knowledge about the context and properties of measuring instruments that yielded the data. The model of experiment for Millikan's oil-drop experiment for measuring the charge of the electron included Newton's laws and the laws of electrostatics. The application of a model of the experiment applied to the data then results in a "data model." For example, the model of experiment may direct us to interpret a sample mean as an estimate of the mean in the population from which data were drawn. A data model might represent the data as a whole, as, for example, by representing data points with a curve, or the data model might represent specific features of the data, such as a sample mean or value of charge of the electron.

Data models do not always arise from the digestion of statistical data or other manipulations of noted-in-the-lab-book observations. Observations are often made with sophisticated instruments, such as an electron microscope or a bubble chamber, that issue in concrete objects, such as a micrograph or a picture of bubble chamber tracks. These objects must be interpreted, again using a great deal of often highly theoretical information deployed in a model of the specific instrument or experiment. These interpreted concrete objects can also be thought of as data models inasmuch as the theoretical interpretation applied from the model of experiment yields something much more richly representational than the uninterpreted micrograph or bubble chamber photograph, or the like, as they might be viewed by someone with no relevant technical expertise (Harris 2003).

https://doi.org/10.1086/656537 Published online by Cambridge University Press

Details will vary greatly from subject to subject, topic to topic, and from one specific observational or experimental arrangement to another. What is common is an output, a data model, concrete or abstract, that carries with it a great deal of interpretation that goes far beyond what might be said, in a naive sense, to have been observed. It is such data models, not pretheoretical objects of observation, that one then compares with a specific hypothesis or generalization for fit or lack of fit, either as a whole or in some more specific respect.

I have sketched a view of theory and data as viewed from the first decade of the twenty-first century. Where in all of this are B&W's phenomena from 2 decades earlier? Remember, B&W's primary objective was to point out that, generally, theory is not compared directly with raw, uninterpreted data. Their conception of phenomena was an analytic tool designed to make this point clear. Whether put in terms of "phenomena" or in terms of a "hierarchy of models," there is a lot of processing that goes on between abstract theoretical principles and raw data. At this level of description, the contemporary hierarchy of models bears out B&W as wonderfully as could be desired. It makes exactly the same point.

With the caveat that all these ideas are flexible and will vary quite a bit from case to case, I think that we can also see a more detailed coincidence between the older and newer approaches. We can take data models to be representations of phenomena tokens and models of experiment to determine what generalizations/phenomena types these should be taken to be tokens of. The phenomena types are the relatively low-level hypotheses and generalizations or particular aspects of these. The data models/phenomena tokens provide evidence for concluding that various low-level generalizations in the top half of the hierarchy provide good, though not perfect, representations of their targets, and through intermediate-level models, for the fruitfulness of the model-building principles at the highest level. In the other direction, highest-level principles function through the offices of intermediate layers to explain relatively low-level hypotheses and generalizations and, through these, to explain the phenomena tokens represented by individual data models. Idealization smooths virtually every step downward as well as upward in the hierarchy. In these ways the hierarchy of models illuminates the connections between idealization and phenomena types and tokens.

The modeler's hierarchy differs from B&W's presentation of phenomena in adding additional structure. I see no important conflict. "Saving the phenomena" presented a prescient way of thinking about the relation between data and theory that has been usefully elaborated—just what is to be expected of any good idea.

https://doi.org/10.1086/656537 Published online by Cambridge University Press

REFERENCES

Bogen, James, and James Woodward. 1988. "Saving the Phenomena." *Philosophical Review* 97:303–52.
Brading, Katherine. 2010. "Autonomous Patterns and Scientific Realism."*Philosophy of Science*, in this issue.
Duhem, Pierre. 1969. *To Save the Phenomena*. Chicago: University of Chicago Press.
Giere, Ronald. 2010. "An Agent-Based Conception of Models and Scientific Representation." *Synthese* 172 (2): 269–81.
Gooday, Graeme. 1995. "The Morals of Energy Metering." In *The Values of Precision*, ed. Norton Wise. Princeton, NJ: Princeton University Press.
Harris, Todd. 2003. "Data Models and Acquisition and Manipulation of Data." *Philosophy of Science* 70:1508–17.
McAllister, James. 2007. "Model Selection and the Multiplicity of Patterns in Empirical Data." *Philosophy of Science* 74:884–94.
Teller, Paul. 2004. "The Law Idealization." *Philosophy of Science* 71:730–41.
———. 2008. "Representation in Science." In *The Routledge Companion to the Philosophy of Science*, ed. Stathis Psillos and Martin Curd, 435–41. New York: Routledge.
Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.
———. 2007. "Causation with a Human Face." In *Causation, Physics and the Constitution of Reality*, ed. Huw Price and Richard Corry. Oxford: Oxford University Press.

https://doi.org/10.1086/656537 Published online by Cambridge University Press