

To appear in a possibly expurgated version in *Metaphysics and the Philosophy of Science* (ed. Matthew Slater and Zanja Yudell)

Interventionism and the Missing Metaphysics: A Dialogue

Jim Woodward
HPS, U of Pittsburgh

Apologetic Preface:

A number of philosophers with a metaphysical orientation have criticized *Making Things Happen* for its failure to provide an account of the metaphysical foundations or grounds or truth-makers for causal and explanatory claims. I originally attempted to write an ordinary paper responding to these criticisms but found this to be a very difficult undertaking: I realized that I disagreed with my critics about so much that putting everything into an ordinary “linear” argument was impossible. I also realized that important elements in my disagreement had to do not just with the ideas of the critics but with the rhetorical devices and strategies for conversational control that they sometimes employ. This led me to the idea of writing a dialog that attempts to highlight the latter as well as a former. What follows is, I readily admit, a caricature which makes no attempt to be fair or balanced. Many of the philosophers I know who are analytical metaphysicians do not share the affect and attitudes of my Professor Metafisico and many are far more knowledgeable about science. I hope that readers will take the dialog in the spirit in which it is intended—as an attempt to be provocative and to raise in sharp, unnuanced way some questions that deserve more attention than they have hitherto received. These include issues about just what metaphysical grounding consists in or amounts to, why (or when or for what purposes) it is required, and how providing metaphysical foundations relates to providing scientific explanations of a more ordinary empirical sort and to methodological concerns that at least in the past were regarded as an important component of philosophy of science. In particular, I want to raise the question of whether it is somehow required that all philosophers of science do metaphysics or (as I maintain and hope) there are kinds of inquiry in philosophy of science having to do with methodology and the interpretation of the content of the particular scientific theories that can be pursued independently of the kinds of concerns that animate analytical metaphysicians.

Other themes that I have tried to illustrate include the following:

I am grateful to Bob Batterman, Jim Bogen, Craig Callender, Chis Hitchcock, Peter Machamer, John Norton, Matthew Slater, Kyle Stanford, Martin Thomson-Jones and Mark Wilson for comments on an earlier draft. In this case, the usual disclaimer that they do not necessarily endorse the ideas expressed in this dialog applies with extra force.

1. Ambivalence about science: Many contemporary metaphysicians express great respect for science and even claim that what they are doing is in various ways continuous with science (either in content or method or both). Yet at the same time they make substantive claims and employ methods that don't seem to be part of any scientific discipline, as in their frequent appeals to "intuition" or to notions like metaphysical explanation¹. This willingness to work both sides of the fence is a great argumentative resource, but a bit frustrating to interlocutors.

2. Conversation Control. One of the most striking features of contemporary metaphysical discussion, at least to an outsider like me, is the use of rhetorical strategies or framing assumptions that both automatically render certain considerations irrelevant and automatically render other considerations central to the discussion, regardless of what you, the non-metaphysician, would like to be talking about. In effect, these devices are employed to *force* the conversation to be only about issues the metaphysician regards as important. One device for doing this is the adoption of highly expansive conceptions of "metaphysics" so that if, e.g., you try to define anything or clarify a concept you are making a "metaphysical" claim. Similarly, ordinary empirical claims of some generality (e.g. about the dimensionality of spacetime) are taken to be metaphysical claims². Given that virtually everything is metaphysics (or at least that boundaries of "metaphysics" are very elastic and ill-defined), it becomes difficult to argue that one is not doing metaphysics and the non-metaphysician becomes vulnerable to Metafisico's rhetorical maneuvers (since you are already doing metaphysics, you might as well do it right, which requires a discussion of grounding or something similar.) Another device is the invocation of various distinctions such as a contrast between those features of science that have "merely" epistemic or methodological significance and those that have true ontological/metaphysical significance, with only the latter being treated as a fit subject for philosophical discussion. The language metaphysicians use to characterize their concerns (frequent use of words like "deep", "fundamental", "ontologically serious" and so on, with the apparent implication that non-metaphysical projects are superficial, shallow and non-serious) has a similar rhetorical role in directing philosophical discussion, as does the use of phrases like "merely pragmatic" as a term of denigration. I will add that I'm convinced that most metaphysicians are completely unaware that the rhetorical devices they employ look to those who are not metaphysicians as though they are tailor made to achieve such exclusionary effects—instead they think of the devices as both natural and completely innocent.

¹ In addition, I think that a great deal of analytical metaphysics is animated by mistaken ideas about the methods actually employed in successful science. However, this is a topic for another paper.

² To clarify: I'm not opposed to all expansive conceptions of metaphysics. If one engages in, say, some wide-ranging, empirically informed exploration of human cognition and behavior and calls that "metaphysics", far be it from me to object. What I object to is the use of such an expansive conception as a device to justify the characteristic preoccupations of analytic metaphysicians in the manner illustrated below.

3. Preferences and Arguments. Related, to (2), expressions of subjective preferences for (or value judgments regarding) certain kinds of philosophical work or topics presented as though they are “objective” arguments that require that one only work on these topics. From my point of view, it would be very refreshing if more metaphysicians would just say that they find doing the metaphysics of grounding more interesting or worthwhile than, say, trying to understand the role of experimentation in causal inference and that for this reason they prefer to work on the former. Instead, what we too often get are claims that work on the former is logically required if one works on the latter.

4) Details: If you think that you already know that some philosophical position of a certain general type *must* (for metaphysical reasons) be correct, you may not be too concerned about working out the details of that position. Apparent problems can be dismissed because (since your position is correct) you know it must be possible in principle to meet those objections.

5) Methodology. There are many issues and problems, arising in the various sciences, that are, broadly speaking, epistemological or methodological in nature. These have long been regarded as legitimate and important topics for philosophical discussion and, as I see it, are central to philosophy of science. They include, for example, investigation (both descriptive and normative) of different strategies for learning about nature (including investigations of the limits of what can be learned from various sorts of evidence), and the characterization and evaluation of various forms of reasoning that figure in different areas of scientific and non-scientific investigation. Work of this sort is pursued not only in philosophy of science, but also in statistics, machine learning, discussions of experimental design, artificial intelligence, and portions of cognitive psychology, not to mention the various sciences themselves. I would also include under this general heading of methodology efforts at clarification of concepts that are central of the various particular sciences, as in Kendler (2005). Analytical metaphysics and allied work in the metaphysics of science seems to me to take little notice of methodological issues and concerns of the sort just described. This is partly a consequence of the tendency of contemporary metaphysicians to downplay or attempt to abstract away from epistemic and “pragmatic” considerations, which makes it difficult to discuss anything relevant to methodology. What follows is in part an attempt to push back against this tendency and to argue for the independent value of methodological concerns.

The Dialog

JW: Hi. I’m Jim Woodward. Pleased to meet you. And your name?

M. I am Professor Metafisico, Ph.D.

JW: What do you work on?

M: I'm an Analytical Metaphysician.

JW: I'm a philosopher of science who works on causation, among other things.

M: That sounds interesting. I'd like to hear the details.

JW: I try to develop the idea that causal claims describe relationships that are potentially exploitable for purposes of manipulation and control. Put slightly differently causal claims describe the outcome of hypothetical experiments: very roughly, X causes Y if and only if there is some intervention I that changes the value of X such that if I were to occur, the value of Y would change. An intervention on X with respect to Y can be thought of as an idealized, unconfounded experimental manipulation of X . When the relationship between X and Y is a causal relationship, it will be stable or continue to hold—*invariant*-- under at least some interventions on X . Causal relationships differ in their degree of invariance, with some causal relationships being invariant under a wider range of interventions and other sorts of changes than others. There is more detail in my book, *Making Things Happen (MTH)*.

M: Hold it right there. I hate to break the news to you, but your idea is a non-starter -- it sounds like you're saying that whether X causes Y *metaphysically depend* on what would happen to Y under interventions on X . But that can't be right. The claim on the right hand side of your definition is a counterfactual. Counterfactuals can't be "barely true"—they require "truth-makers" or "grounding"³ in terms of what is non-counterfactual (i.e. categorical) and more fundamental. "Truth supervenes on Being", as we metaphysicians like to say. Interventions and even more so claims about what would happen under interventions are not plausible candidates for what is Ontologically Fundamental.

JW: Actually, the above "definition" or bi-conditional was not intended as a claim about metaphysical dependence or anything like that. I put it forward because I thought it captured the way that causal claims are used or understood in a number of areas of science—especially but not exclusively the social behavioral and biological sciences—and in common sense reasoning about causation and because I thought this in turn had interesting methodological implications. I claimed in *MTH* that this way of thinking about causal claims was fruitful because it helped to clarify what researchers in the above disciplines mean or are committed to (or are trying to establish) when they make causal claims and also how such claims relate to ideas about causation advanced by philosophers and researchers in other disciplines.

³ There are a number of different positions and accompanying terminology that are in the general neighborhood of Metafisico's views. Some writers talk about "grounding" causal and other sorts of claims, others about identifying the "truth-makers" for such claims or about specifying what such claims "metaphysically (or ontologically) depend on". While acknowledging these differences, in an effort to make this dialog readable I have deliberately lumped these different positions together. So for "grounds" the reader should read: "grounds or truth conditions or...".

I also suggested that interpreting causal claims in this way had a normative or regulative significance: Vague or unclear causal claims may be made disambiguated by spelling out the hypothetical experiments with which they are associated. A related point is that we can often clarify the sort of evidence and other assumptions needed to establish a causal claim on the basis of non-experimental data by reflecting that this must be evidence that is sufficient to establish what the outcome of the associated experiment would be, were able to perform it. This helps us to see why and when certain techniques such as the use of instrumental variables in econometrics can be used to reliably establish causal conclusions. I've always thought of these goals as interpretive/descriptive/semantic/methodological (methodological for short) rather than metaphysical-- somewhat similar in aspiration to the treatment of causation in Holland (1986) or in King, Keohane, and Verba (1994) or Morgan and Winship (2007) and having some affinities with and drawing on ideas from Spirtes, Glymour and Scheines, (2009) and Pearl (2009). Many of these writers also offer what they call "definitions" of various causal notions, but this is more in the service of the methodological goals I mentioned and not because they are interested in making claims about "metaphysical dependence". In fact, definitions abound in science and mathematics but they are rarely if ever intended as claims about metaphysical dependence—books on real analysis typically define the continuity of a function in terms of the notion of an open neighborhood, but I don't think they mean to claim that continuity metaphysically depends on this notion. I'll add, though, that from my point of view nothing really turns on my use of the word "definition"—I'm happy to drop it in favor of talk of relationships of mutual constraint between causal claims and claims about the outcomes of hypothetical experiments or in favor of talk about using the latter to elucidate or clarify the former with an eye to the goals described above⁴.

M: These people you describe above—I haven't heard of any of them. Where did they go to philosophy graduate school? Who did they study Metaphysics with?

JW: They don't have Ph. D's in philosophy and they aren't aiming to do metaphysics. Holland is a statistician, King a political scientist, and Pearl a computer scientist.

X: Well, I don't mean to be dismissive but it seems unlikely in that case that they will be able to make any contributions to the Fundamental Ontology of causation. It sounds like what they are doing is at best of merely practical or pragmatic interest.

JW: I'm OK with that. More than OK in fact. I think it is worthwhile to do things that are of merely pragmatic interest. That is precisely my point. They aren't trying to make contributions to Metaphysics. They are trying to do something else and similarly for me.

M: Well, the work you describe may be of interest to some people outside of philosophy in its unambitious way, but it doesn't sound as though what either you or they are

⁴ Some of the ideas described in this paragraph are defended in more detail in Woodward, forthcoming.

proposing has any “metaphysical depth” to it⁵. I fear that it lacks what C.B. Martin and others following him have called “ontological seriousness”.

JW: Once again, I’m happy to agree. It seems to me that the people I mention just have a different set of interests from you. And the same is true for me. Why can’t we just agree on a division of labor? You can do metaphysics and I’ll do what interests me.

M: What you call a division of labor is not possible. By *definition* metaphysics is concerned with “the foundations of reality” (Chalmers, Manley and Wasserman, 2009, p.1), with “Being as such” (van Inwagen, 2007), and with “the fundamental structure of the world” (Sider, 2009, p. 420) Everything is thus suffused with or at least grounded in metaphysics and, as a result, you can’t *not* do metaphysics (Sider, p. 420) . Indeed, if you deny you are doing metaphysics, you are by that very denial doing metaphysics since you must have in mind some conception of what it is that you are not doing and you must think that it is possible not to do that. This itself amounts to a metaphysical position which needs to be defended—by doing metaphysics of course. Oh, and I almost forgot, if you discuss the ideas of anyone who is a metaphysician you also must be doing metaphysics (Strevens, 2007) .

JW: It sounds as you are saying that I have to do what you are doing, according to your rules. And you don’t have to provide any motivation or justification for what you are doing, since one can’t not do it.

M: Your description of the dialectical situation is essentially correct, although rather unsympathetically expressed. Let me emphasize, though, that these requirements concerning the centrality and inescapability of metaphysics do not in any way reflect the personal ambitions of us metaphysicians. Instead, these requirements flow ineluctably from the foundational role of metaphysics in all inquiry. And since, as we have established, you can’t avoid doing metaphysics, you might as well buckle down and do it thoroughly and correctly, even if that means that you never get around to doing any of this stuff you call methodology (aside: which wouldn’t be a bad thing at all, in my opinion).

JW: I would find it easier to understand what you are getting at if you would stop using words like “depth” and “serious” and explain to me exactly what is wrong with being interested in the merely methodological. For example, will I make mistakes in the methodological projects that interest me if I fail to provide grounds/truth-makers in your sense for causal and counterfactual claims? Will this lead to empirical mistakes of some kind?

M: How can any area of inquiry be satisfactory if it is left ungrounded and without foundations? As for specific problems that arise for your project if you fail to provide metaphysical foundations, I haven’t worked out the details since I don’t think this is

⁵ Approximately quoting a prominent philosopher of mind commenting on my views at a recent conference. He didn’t intend this as a compliment, but I was happy to agree.

necessary, but I think it is obvious that you'll end up with an account that is unclear, incomplete, and subject to other problems as well. And of course it won't be Ontologically Serious and won't be a contribution to anything that is Fundamental.

JW: But the people mentioned above don't provide metaphysical foundations and yet what they say doesn't *seem* to be confused. Or if what they say is defective in some way, this doesn't seem to be because they have failed to provide metaphysical foundations. Can you point to specific problems that arise with their accounts because they fail to provide metaphysical foundations? For example, Rubin (1974) and Holland (1986) employ what is in effect a counterfactual treatment of causation, defining the causal effect of a treatment t with respect to an alternative treatment t' for a unit u as $Y_t(u) - Y_{t'}(u)$ where $Y_t(u)$ is the value Y would have assumed for unit u if it had been assigned treatment t and $Y_{t'}(u)$ is the value Y would have assumed had u instead been assigned treatment t' . This is then used to characterize various sorts of assumptions that may be used for reliable inference to causal conclusions from statistical data—the idea is that the assumptions must be such that in conjunction with the statistical data, they allow you to estimate $Y_t(u) - Y_{t'}(u)$ or perhaps the expectation of this quantity $E[(Y_t(u) - Y_{t'}(u))]$. These ideas are criticized by some, but these criticisms don't seem to have much to do with any failure on Holland's part to provide a proper *metaphysics* of causation. Pearl, for example, argues that it is often easier and more straightforward (and we are less likely to make mistakes) if we employ graphical methods rather than variables describing counterfactual outcomes like $Y_t(u)$ but Pearl's graphical representations don't supply truth-makers in the sense that you seem to have in mind either and in any case Pearl doesn't seem to be claiming that they are metaphysically superior to treatments like Holland's. My own view is that the notion of an intervention provides a more precise clarification of Rubin and Holland's notion of the “assignment” of a treatment to a unit – a notion which is left underspecified in their accounts -- but again this doesn't seem to turn on issues having to do with metaphysical foundations.

M: It does sound from your description as though the interests of these people are, sadly, merely pragmatic. But exactly because of this, I don't see that what they are doing is of any Metaphysical Interest. You still haven't told me anything about what the truth-makers or grounds are for causal claims.

JW: I thought that we were talking about your claim that various problems will arise for what they (and I) are doing if we fail to provide metaphysical foundations. You seem to be confusing this claim with the assumption that they ought to be doing metaphysics and/or that only metaphysics is worth doing. This assumption amounts to a *preference* on your part for certain kinds of inquiry. You seem to be confusing this preference with an *argument* that what you advocate doing is required or unavoidable.

M: Let me give it another try⁶. A while back you asserted a bi-conditional linking causal claims to the outcomes of hypothetical experiments. And you also described what you were doing as an attempt to capture what those who use causal claims in some areas of

⁶ Thanks to Martin Thomson-Jones for raising some issues connected to what follows.

science “mean” or are “committed to” in connection with those claims and that causal claims can be clarified by associating them with hypothetical experiments. But everyone agrees that when you assert such a bi-conditional, and particularly when you connect it to notions like “meaning”, you are making a metaphysical assertion – you are making a claim about the grounds or truth-makers (or something similar) for causal claims. So I don’t understand how you can claim that you are not trying to do metaphysics.

JW: I don’t understand why the claims that I make about causation *must* be interpreted in the way you describe. You seem to be treating this as a forced move of some kind that I am required to make. As I said above, I’m not particularly wedded to describing what I’m doing in terms of providing definitions or meanings, especially if you are going to then go on to interpret what I’m saying in the metaphysical way you want to foist on me. Perhaps it would be better (at least for the purposes of this conversation) if I put what I am saying in terms of characterizing a certain set of practices for reasoning about causal claims, connecting them up with other sorts of claims, making distinctions among such claims and so on. Describing and elucidating aspects of the behavior of certain concepts (like “cause” and “intervention”) and connecting them to methodological goals, or something along those lines. Surely there are activities of this sort that go on all the time in science and among many who comment on science (both philosophers of science and others), but which we don’t have to view as bound up with metaphysics in the way you have in mind. You seem to be simply stipulating a notion for “meaning” and, for that matter, “if and only if” that builds in the idea that one can’t do what I am trying to do without making heavy-duty metaphysical claims about grounding. Why can’t I just decline to accept your stipulation?

That said, I wonder if it wouldn’t be more fruitful to try another tack. Perhaps you might tell me more about this grounding or truth-making relationship that I’m not adequately dealing with. That way I would have a better sense of what I am leaving out.

M: I’d be delighted to—this is the subject of a rich and flourishing literature (see, for example, Chalmers, Manley, and Wasserman, 2009, and Correia and Schnieder, 2012). Let me begin with grounding, although I should acknowledge at the outset that there is a fair amount of disagreement about just what this relationship involves or even whether it is intelligible and about whether it is or is not the same thing as truth-making. For starters, many think that the grounding relation is primitive, and thus it can’t really be explained in terms of anything else (Schaffer, 2009). But I can tell you about some of its properties. For example, many metaphysicians hold that it is irreflexive, asymmetric and transitive. Its asymmetry distinguishes it from more ordinary sorts of definitional relationships, which may permit you some choice about which terms one starts off with and which terms are defined in terms of them, as is often the case with various alternative sets of definitions and axiomatizations in logic and mathematics. (Think of various choices of primitives in propositional logic or in Euclidean geometry.) This contrasts with proper metaphysical definitions which express relations of metaphysical dependence. Here there is a uniquely correct order of definition: p is defined in terms of its grounds or truth conditions q , q is defined in terms of its grounding conditions and so on until we reach those entities and relations which are fundamental. Grounding

relations thus give us insight into the hierarchical structure of reality, according to which some entities or facts are dependent for their existence on others, while other entities or facts are fundamental, requiring no further grounding. Thus, just as Philosophy programs should be ranked in a hierarchy, according to their Excellence in Metaphysics, so also Reality itself can be ranked, in terms of its relation to what is Fundamental. To be is to be Ranked⁷.

JW: Fascinating. Tell me more.

M: The grounding relation has other properties as well. For example, it is hyperintensional, so that logically equivalent claims can have different grounds. The grounding relationship is also related to the best theory we have of everything—the best theory should be thought of as a theory of what is fundamental that correctly tracks grounding relations (Schaffer, 2009). Among scientifically minded philosophers, it is assumed that the best candidate for such a theory is to be found in Fundamental Physics—a view to which I myself subscribe. So the grounding entities (or whatever) will presumably include, for example, elementary particles and fields of the sort discussed in fundamental physics. Fortunately, though, we needn't inquire too closely into the details of this—we can take this aspect of what is fundamental just to have to do with “point particles and fundamental physical magnitudes” (Schaffer, 2009, P. 370). So to sum up, when we metaphysicians ask for the grounds or truth makers for causal claims or claims about laws what we are asking for are facts or entities that stand in the kind of relationship to causal claims and laws that is common to all of the examples described above.

JW: I'm still a little confused. On the one hand, it sounds as though metaphysics and grounding, as you describe them, are (at least among scientifically minded metaphysicians) closely connected to physics—apparently to what you think of as fundamental physics. On the other hand, I take it that metaphysics is not exactly the same thing as physics? So what exactly is their relation?

M: Obviously metaphysics is not the same thing as physics. If it were, what would be left for metaphysicians to do? The very fact that metaphysicians keep doing metaphysics, as they have done for thousands of years, and do it, at present, in a way that has no influence whatsoever on the practice of physics, shows that what we are doing must be importantly distinct from physics and the other sciences. As for the relation between these disciplines, this is a matter of on-going discussion. One possible view is that metaphysics is inspired in part by physics, but must satisfy other constraints too -- for example, consistency with the intuitions of metaphysicians about grounding. Physicists, although admirable in many ways, tend to be inattentive to the need for grounds. For example, they may introduce claims about laws of nature, and support these only with empirical evidence, and physical arguments about symmetries and so on, while failing to specify what the ultimate grounds for these law-claims are. One useful thing that

⁷ Or, in a more “continental” mode, Being Ranks Itself.

metaphysicians can do is to step in and provide more adequate grounds for these claims. It is your failure to do this that makes your views about causation so defective.

Another possible view of the relation between science and metaphysics (perhaps not sharply distinct from the view just described) is that metaphysicians provide accounts of the basic categories in terms of which all scientific theories must be framed—categories like cause, law, substance, part and whole and so on. Scientists then fill in the details, providing particular examples of causal relationships, laws and so on. So in this picture as well, metaphysics plays a foundational role in science.

Yet another other thing that distinguishes us metaphysicians from what physicists do is that we are not so interested in, so to speak, actually existing physics, but rather in the wider space of metaphysically possible worlds. Our aim is to discover what the grounds would have to be for any metaphysically possible physics⁸. This is why we are concerned to make our theories physics-like, or physics-inspired rather than being too closely tied to the physical details of our world, since the latter course would lead to a regrettable loss of generality. For example, in discussing laws of nature, we find it most useful to think of these just as generalizations of the form “All As are Bs” and in discussing scientific theories we represent them as theory *T*. Any further detail would just be limiting.

Perhaps I should add that as a scientifically oriented metaphysician, I fully accept Quine’s dictum that metaphysics is continuous with and “on a par” with science and that metaphysicians are just doing at a somewhat more abstract and general level what those physicists at CERN who are looking for the Higgs boson are doing. The criteria for a good metaphysical theory are at a general level the same as those for a good scientific theory, simplicity, great scope, theoretical fertility and so on, as many metaphysicians have said (e.g., Sider 2009). On the other hand, it also needs to be borne in mind that we are doing something that is more fundamental than any particular science and that shouldn’t be judged just by the sorts of pragmatic criteria that are unfortunately sometimes employed in the assessment of scientific theories, like usefulness in prediction and control, computational tractability and so on.

A good way of summing up is that what we are doing is both science-like and not science-like. I hope you find this clarifying.

JW: It would help me if you could provide some more specific examples of grounding.

M: I’m happy to oblige. The literature is full of illustrations: the disjuncts of a disjunction ground the disjunction, geology is grounded in physics, according to some philosophers causal claims are grounded in facts about dispositions, according to other philosophers causal claims are grounded in facts about laws of nature, and laws of nature are grounded in relations of necessitation between universals. According to still other philosophers, laws of nature are instead grounded in facts about the deductive systemization of the entire Humean Basis that achieves a best balance of simplicity and strength. Other illustrations: fists (if they exist—there is dispute among ontologists about this) are grounded in hands or shapes assumed by hands, and the set whose sole member

⁸ Cf. Lowe (2011) who characterizes metaphysics “as *the science of the possible*, charged with charting the domain of objective or real possibility...” (p.100)

is the empty set is grounded by the empty set. However, I caution that these are just examples. The grounding relation is the single unique relation that is common to all of these particular cases. Thus when we metaphysicians ask for the grounds or truth makers for causal claims or claims about laws what we are asking for are facts or entities that stand in the relationship to causal claims and laws that is exemplified in all of the above examples.

JW: I'm afraid I'm still not getting it. Why do you think that these are all common exemplifications of a single grounding relationship? It looks to me that in the first example, what you have is an asymmetry of logical entailment between a disjunction and its disjuncts. In the second case, you are appealing to some idea according to which the entities that are the subject matter of geology are made up of entities like atoms that are studied in physics. So in this case you are talking about some sort of compositional relation. Or perhaps you are claiming the geological phenomena can be explained in terms of (or "reduced to") relationships that figure in physics where "explanation" here means scientific or physical explanation. In still other cases, your grounding relations seem to be conceptual or definitional in character. At any event, when you describe all of these cases in terms of grounding or truth-making, don't you risk conflating causal or explanatory dependence (of the sort discovered in empirical science) with logical entailment with part/whole relations with various sorts of semantic or definitional relations with...? I've always thought that one of the great achievements of early modern era in philosophy was the clear separation of these notions, which tended to be conflated in earlier scholastic thought. Indeed, don't locutions like "truth-maker" encourage such conflation, since "making" sounds like a causal notion, but at the same time truth-making relations seem to be conceptualized as something very different from ordinary causal relationships? Perhaps you should consider the possibility that there is no single relationship that has all of the features of grounding or truth-making, at least as you are conceiving of it.

M: I don't see how that is possible. Grounding and truth-making have been central to metaphysics since Aristotle, if not Thales, and discussion of these continues in the Best Departments today. How could all of these people be wrong about whether there is such a relationship? And if there is a grounding relationship, there could hardly be more than one, could there? I mean, what would happen to Fundamental Ontology (not to speak of our Departmental Ranking) then? I would say that the very fact that metaphysicians talk (and disagree) about grounding (and/or truth-making) shows that these notions (or at least one of them) must be intelligible and that is our job to discover its properties.

JW: I see. Let me then try to express a somewhat different worry. You agree that there is considerable disagreement among metaphysicians about what grounding (or for that matter related notions) involve and how to understand this notion.

M: Oh, yes. This is a very lively area of debate, with a wide spectrum of different views.

JW: But then perhaps you can also see that this puts me at a considerable rhetorical disadvantage when metaphysicians like you demand grounds for causal claims and

counterfactuals. I mean, if you cannot agree among yourselves about what you are talking about when you ask about grounds, it is going to be difficult for someone like me to respond to your demands.

M: That is your problem, not ours. I can tell you that it probably will take some time for us metaphysicians to work out the full theory of grounding. And perhaps, alternatively, we will decide that talk of grounding should be replaced by something else (Truth-making, Structure or what have you). But we are all agreed that grounding (or something like it) is crucially important and it is a great failure on your part that you have not provided a detailed story about how this works in connection with causal claims and counterfactuals, even if we have not yet fully decided what grounding is.

JW: Perhaps we are more likely to get somewhere if we talk more about specific examples. Consider the behavior of a macroscopic sample of dilute gas conforming to the ideal gas law and other familiar generalizations of thermodynamics. In this case would the underlying “grounds” for this behavior consist in, say, facts about its constituent molecules and the laws governing their interactions of the sort that would be described in statistical mechanics?

M. Something like that. As I’ve already suggested, we metaphysicians largely leave the details of which grounding relationships are actually exemplified in our world to physicists, except when they make obvious metaphysical mistakes. Presumably the true grounds in the example you describe involve far more fundamental physics—truths of string theory or something similar. In any case, the details don’t matter—it is the general idea of grounding that is philosophically important. This is why merely providing this underlying physics does not by itself answer the metaphysician’s concerns about the proper elucidation of the grounding relationship.

JW: Well, here is something that puzzles me. Many, many different assumptions about the micro-constituents of matter and their interactions are consistent with the holding of the ideal gas law—even the assumption that the gas is a continuum is consistent with this law, given the right general assumptions about the behavior of this continuum. The same is true of many other aspects of the behavior of macroscopic objects—their behavior is surprisingly independent of the details of the underlying physics, often depending only on very generic features of that physics. For example, aspects of the behavior of many substances of quite different material composition undergoing phase transitions depends only on very general features of those systems—their dimensionality, the symmetry properties of the Hamiltonians governing these systems and so on. This is why, to paraphrase Nigel Goldenfeld and Leo Kadanoff (1999), to model a bull-dozer you don’t have to model its constituent quarks. So my question is this: why aren’t the grounds or truth conditions (or whatever) for this behavior these more generic features or perhaps a disjunction or equivalence class of all theories consistent with the macroscopic behavior, rather than whatever highly specific “low-level” physical truths about strings or whatever hold in the actual world? After all, in one obvious sense the macroscopic behavior does not depend on these very specific truths at the fundamental level. So if the grounds for the behavior of macroscopic entities have to do, among other things, with

those features to which we appeal to provide ordinary scientific explanations of that behavior, it seems to me that there is a good case to be made that these features in many cases are not going to involve highly specific truths of some fundamental physical theory. Again, perhaps this is a reflection of the fact that you are trying to do too many things at once with your grounding/truth-making idea—these are supposed to reflect underlying physics, provide explanations, provide reductions to supposedly conceptually more primitive facts and much more, but maybe a single relationship can't embody all of these features.

M: I haven't really thought about the issues you describe about the independence of the behavior of macroscopic objects from the physics that underlies them and, frankly they involve details of a sort that bore me. But to the extent you are suggesting that the grounds in such cases are generic or disjunctive facts, your proposal is a non-starter. Really, you seem to have no ear for metaphysics at all. I've already emphasized that disjunctions cannot be metaphysically fundamental but must rather be grounded in their disjuncts. Moreover, it would be contrary to all Metaphysical Intuition if something as abstract and generic as the symmetry properties of a Hamiltonian could serve as a ground. In general, the difficulty of modeling a bulldozer by means of quantum chromodynamics is merely pragmatic and of no philosophical significance. When we put modeling considerations aside, we see that the truth-makers for the behavior of bulldozers indeed are to be found in the behavior of its constituent quarks (or something like that)

JW: Part of the reason why these issues interests me is that it seems to me that one of the great advantages of experimentation, epistemically and methodologically speaking, is that one can often use it to reliably establish conclusions about causal relationships independently of underlying details that would be provided by some lower level theory. The interventionist account attempts to capture this observation. For example, a researcher may be able to establish that some drug provides a cure for an illness without knowing the underlying chemistry of the drug, much less how its action would be modeled in quantum field theory. This is another case of what I was describing above: a kind of independence or decoupling (in this case epistemic independence) of more macroscopic behavior from underlying physical details. From the point of view of methodology and discovery the most reliable sources of knowledge about the behavior of more macroscopic systems often is *not* found in information about what you are calling grounds or truth makers, insofar as these have to do with fundamental physics. Perhaps that is part of the reason I don't find a concern with grounding, even if intelligible, particularly relevant to my interests.

M: I regard these observations as completely irrelevant. Indeed, I'm shocked that you would allow your philosophical views about causation or anything else to be influenced in this way by merely epistemic considerations. The whole point of our discipline is to get behind the merely practical and epistemic and to limn the True Nature of Reality.

JW: OK. Let me try something different. Consider the counterfactual claim that if right now I were to release this cocktail glass I am now holding in my hand, it would fall to

the floor. If you want grounds or truth-makers for this claim, why can't I just say that these have to do with whatever would figure in a deeper physical explanation of why the counterfactual claim is true — e.g., there is the gravitational force exerted on the glass by the earth in accordance with the inverse square law, the fact that as it falls the resisting force of the surrounding air, modeled by Stokes law, will be negligible, for an object of the dimensions of the glass, the initial conditions of the system once the glass is released and...

M: I'm afraid that once again you're completely missing the point. What you have just described is a mere physical explanation⁹ and just a particular example of one at that. To repeat: Metaphysicians are after something different and deeper — the metaphysics in virtue of which your counterfactual is true. You can supply all the physics you want and you still won't have given me the metaphysics that is so lamentably missing from your account. What we are looking for are the metaphysical facts and relationships that underlie and make true all lawful or causal relationships rather than specific examples of such relationships. (At the same time, though, let me remind you, as emphasized above, that what we are doing is continuous with and “on a par with” science). In the case you mention, for example, a metaphysician will want to know what grounds or serves as a truth condition for the inverse square law to which you appealed in your physical explanation. Surely you can see that it is not satisfactory to just take the inverse square law as metaphysically primitive.

JW: I wasn't suggesting that we take the inverse square law as metaphysically primitive, if only because I don't know what “metaphysically primitive” means. I was instead hoping we could find a way of talking about laws that did not get us embroiled in a discussion of what is metaphysically primitive or derivative. But perhaps we could make better progress if you would give me a specific example of what count as a ground or metaphysical truth-maker for a causal or counterfactual or nomological claim. What have metaphysicians proposed along these lines?

M: I'm happy to oblige, although there are lots of candidates, since again this is an area in which there is a lot of disagreement. Some philosophers hold that the truth makers for laws (and presumably also for causal claims and counterfactuals) are to be found in special entities, (or properties, or relationships.) For example, according to Dretske (1977), Armstrong (1983) and Tooley (1987) , the truth-makers or grounds for laws of nature are relations of necessitation between universals, and according to other philosophers laws are made true by facts about dispositions or “active powers” possessed by particular objects. Still other philosophers (e.g., Ellis, 2001, Bird, 2007), though, think that it follows from the fact that objects possess dispositions and other

⁹ To this it might also be added that in many cases, no one has any idea how to connect causal and counterfactual claims in the special sciences to the fundamental laws that “underlie” them. Because of this, it is opaque how to connect these laws to issues about the testing of such claims, or how to reason with them. Observations about the grounding of such claims in fundamental physical law, even if correct, don't seem to help with such issues.

plausible assumptions, that there are no laws of nature, at least as usually conceived. I myself prefer a second kind of approach, the best systems analysis of laws (hereafter the BSA), which strikes me as more scientific in spirit and also promises a satisfactorily Humean reduction of the notion of law.

JW: I'd like to hear more about the BSA. I'm relieved to hear that you prefer it to the other alternatives you describe because frankly, I'm a more than a little unclear about exactly what they accomplish. Suppose I claim that the gravitational inverse square law or the Lorentz force law holds for some domain of investigation. What exactly have I added if I then supplement this with the claim that both are made true by a relation of necessitation between universals? Presumably this necessitation relation must possess features that exactly mimic whatever other features we think are possessed by the Lorentz force law and we don't seem to have any evidence that the necessitation relation holds that is distinct from whatever evidence we have that the Lorentz force law possesses these other features. For example, if we think that the Lorentz force law licenses certain counterfactuals, then we will also postulate that the associated relation of necessitation among universals also operates in such a way that it licenses the same counterfactuals. Or suppose I say that every individual charged object has a disposition to conform to the Lorentz force law and this is what "grounds" the law. Doesn't this (at best) just re-describe the law in some new terminology? Worse, don't we have a sort of illusion of explanation: We claim that the inverse square law or at least the regularity it describes, holds "because" of a relation of necessitation between universals and so on. But I don't see in what sense we have an "explanation" here, rather than (at best) just a redescription.

M: Well, although I prefer the BSA, I think that, speaking as a metaphysician, you are being unfair to the views just described—they needn't be trivial in the way you suggest. For example, some versions of dispositional accounts make the important and distinctive claim that the truth makers are located "in" individual objects in the form of dispositions they possess, as opposed to being located somewhere else that is external to those objects. Stephen Mumford (2004) uses this idea to argue that there are no such things as laws of nature, at least as ordinarily conceived. This is certainly a non-trivial claim that has generated a great deal of discussion. If correct, it shows that most physicists misunderstand their own discipline—a result that we metaphysicians would welcome, since it would demonstrate the great importance of our investigations.

As for your claim that accounts are just re-description, I agree that they do not provide ordinary scientific or causal explanations. Instead they aim to provide a special kind of explanation— a "metaphysical explanation".

JW: I'm guessing that we wouldn't find it very profitable to discuss why what you call "metaphysical explanations" should be regarded as explanatory. I'll add that I would be inclined to regard it as a liability rather than a virtue of the search for truth makers that it apparently can be used to generate questions about where the truth makers for laws of nature are "located". But since we both agree that the special entities accounts are unpromising, why don't we move on the BS approach?

M: Excellent. Here is the basic idea: Start with the full Humean Basis (HB) — a specification everything that actually happens throughout the entire lifetime of the universe, but purged of any references to modality, possibility, causal or lawful relations, or anything like that. Then consider alternative systemizations of this HB, looking for those that achieve a best balance of simplicity and strength. The laws will be those regularities that are described by the axioms and theorems that are common to all such best systemizations. Assessments of simplicity are to be made with reference to a canonical or privileged language which consists of predicates corresponding to the “perfectly natural properties”.

JW: And why exactly should we believe that laws are connected in this way to finding a best balance of simplicity and strength?

M: The BSA has a clear answer to this question: these standards -- simplicity, strength and achieving a best balance between them -- are the standards (in fact, the only standards) that scientists actually use in choosing among competing theories. And when they are applied to problem of laws they yield an account with all sorts of attractive features. For example, the resulting account satisfies our aspiration for a Humean reduction of claims about laws to claims about regularities (or at least to claims about the way that regularities fit together in a total system).

JW: So you are saying that these are the standards that are actually employed in scientific practice in choosing among alternative theories and that when one chooses in accord with them, one picks out exactly those regularities that scientists classify as laws as opposed to those they regard as accidental? These sound very much like broadly empirical claims about scientific practice—very much the sort of descriptive/interpretive investigation that I told you that I was interested in when our conversation started. So maybe we are not so far apart after all. Claims about metaphysics and grounds/truth – makers are at least in part tied to empirical claims about the content and practices of the various sciences, especially physics.

M: Well, sort of, but I have to admit things are a bit more complicated. It turns out for example, that when you apply the BSA to cosmology, various regularities (such as those having to do with large scale isotropy and uniformity of mass-distribution of the universe) that most cosmologists seem to regard as accidental or non-lawful turn out to be laws after all (Cf. Callender, 2004). So if the BSA is correct, we should conclude that these scientists were wrong in classifying these regularities as non-laws.

JW: I can certainly understand why naturalistically minded philosophers of science prefer the BSA to inflationary metaphysical programs that populate the universe with relations of necessitation between universals and their ilk. I applaud the attempt to provide an account of laws that connects with features of scientific practice. But I don't want to accept the BSA simply on the grounds that it appears less metaphysically extravagant than its competitors. And in this connection, what you've just said concerns me a bit. I mean first you motivate the account by appealing to certain standards that you claim guide scientific practice in the identification of laws, and then when those standards

apparently lead to generalizations being identified as laws that are not so regarded in scientific practice, you conclude that the practice is wrong. Why don't you instead conclude that your account of the practice is wrong—maybe laws don't have to do with trade-offs of simplicity and strength in the way that you claim.

M: While it is true that, as I emphasized earlier, many of us like our work in metaphysics to be science-inspired, we certainly don't want to be uncritical slaves to the pronouncements of scientists. Maybe cosmologists are just wrong in regarding generalizations about the large-scale isotropy (and so on) of the universe as non-laws. That is what our best theory of laws of nature suggests. You don't think we should automatically defer to everything scientists say, do you? As heirs to a proud humanistic tradition that stretches back thousands of years, we should be self-confident enough to avoid such "scientism."

JW: It would help me to understand how the BSA is supposed to work if we could work through a simple example.

M. I'm happy to oblige. Consider a record of the positions, velocities and masses of the sun and all of the various planets in the solar system throughout all time. The Humean Basis will include this sort of information. We want to systematize as much of this information as possible within a framework that achieves a best balance of simplicity and strength. Now this record will include the information that the planets (if we include the asteroid Ceres) exhibit a regular pattern or spacing in the distance of their orbits from the sun—this is what is known as Bode's law¹⁰. Bode's law thus describes a regularity. Consider adding this to the Best Systemization as an axiom. The result of doing so would be (let us assume) a gain in strength since we will now be able to use this regularity to deduce more facts about the behavior of the planets. At the same time, addition of Bode's law to the BS would result in a loss of simplicity since the systemization now includes an additional axiom. In assessing this trade-off, we see that this loss in simplicity outweighs the gain in strength, hence that the Bode's "law" does not belong to the Best Systemization and is consequently not a law.

JW: This is ingenious but I'm not convinced that it is a very accurate description of the way in which the epistemic values that guide the discovery of laws relate to one another. For one thing, even if we put aside the fact that we do not have a good account of simplicity of the sort this framework requires, theory construction in science does not seem to be guided by the kind of trade-off between simplicity and strength you describe. Rather than a trade-off, science (and particularly fundamental science) often seems guided by an ordering of these two considerations in which strength has something close to lexical priority over simplicity. Other things being equal, it is permissible, even required, to add complexity, even a lot of complexity, when this contributes to even "small" (assuming we know how to measure this) increases in strength. Consider Newton's first rule of reasoning: "We are to admit no more causes of natural things than

¹⁰ Here I am merely reporting a story that is sometimes used to motivate the BSA. In fact the story is inaccurate since Bode's law does not correctly predict the orbit of Neptune.

such as are both true, and sufficient to explain the appearances'. This rule seems to say that if some additional complexity (in the form of an additional cause) is required to explain appearances, we should admit it, but not otherwise. The rule does not say that we should *omit* to introduce additional causes that are sufficient to explain appearances if the simplicity gain resulting from this omission outweighs the resulting loss in strength which is what the trade-off postulated in the BSA seems to require. Instead, Newton seems to be suggesting that the introduction of additional causes/complexity is always justified (perhaps required) if this leads to an improvement in strength. Your picture of trading off simplicity and strength seems to license arguments like the following: my theory doesn't predict very much about planetary positions or movements but it is so much simpler than theories that do that it achieves a better overall balance of simplicity and strength and is preferable for that reason. Isn't it obvious that this is not supported by generally accepted ideas about good scientific methodology?

Moreover, if this is right, then your story about Bode's law does not work. Instead, on the assumption that strength has (near) lexical priority over simplicity, it seems we should conclude that Bode's law and other similar generalizations should after all be admitted as axioms in the Best Systemization (since their addition to the BS would lead to an increase in strength), in which case they will count as laws, contrary to what everyone seems to agree is the correct assessment of their status¹¹.

M: I agree that the details of the BSA may require a little tweaking—some of us are working on that. The important thing to bear in mind, though, is that simplicity and strength are important in theory-choice and that whatever the details may be, we can use them to provide a properly reductive account of laws. I'd advise you to focus on this larger point and not get too distracted by niggling details. The BSA is the most metaphysically promising approach to laws of nature that we have, particularly since it is appropriately reductive. So it *must* be right, at least in broad outline. Otherwise we would have a situation in which the laws of nature lack metaphysical foundations and are ungrounded, which is obviously intolerable.

JW: Let me try a slightly different angle. You've said that the BSA is to be thought of as a systemization of the entire Humean Basis— everything that happens in the entire universe over its entire life-time.

M: That's right.

JW: But obviously scientists don't have access to any thing like this when they formulate and test theories. So in this respect at least, the BSA does not describe how science actually proceeds.

M: It's a thought experiment. We project from what the small corner of the universe to which we presently have access to everywhere and everywhen and then think in terms of competing systemizations of that.

¹¹ For more detail, see Woodward 2014.

JW: But this Humean Basis we end up imagining—it just involves at bottom what you describe as point-wise arrangements fundamental physical magnitude or something like that? And this is understood in such a way that no characterizations making reference to any kind of physical modality is allowed into the characterization of the basis?

M: That's right. If we let that stuff in we would, we would lose the possibility of a reduction. We would be caught in a circle.

JW: But now it looks as though it is your aspiration to produce a reduction rather than anything grounded in how scientists actually reason about laws that is driving your characterization of this supposed basis. In one respect, the HB includes far more information (about what happens everywhere in the universe over its entire life-time) than scientists in fact have access to. In another respect, the Basis seems to include far less information, in comparison with the evidence and other considerations on which scientists actually rely in reasoning to conclusions about laws, causal relationships or other claims with modal content. What I mean is that when I look at the ways in which scientists in various disciplines, including physics, infer or reason to conclusions with modal content, they often seem to make use of some version of the following schema:

Prior assumptions about modal (causal or nomological) relationships p + non-modal information I about regularities, correlations, initial and boundary conditions \rightarrow conclusions about other modal relationships (new modal knowledge) p^* .

For example, prior causal knowledge (e.g. that Y does not cause X and that X and Y are not caused by some third variable Z) can be combined with information about correlations (e.g., that X and Y are correlated) to infer that X causes Y . Information about the trajectories of the planets of the sort represented by Kepler's laws can be combined with very general nomologically committed assumptions (e.g. the motion of the planets is caused by some central force due to the sun) to derive the gravitational inverse square law, as Newton himself showed.

If your characterization of the HB excludes all such prior causal or nomological information, how can it capture how scientists actually reason when they make causal or nomological inferences? And in general if our aim is to capture "our" (or scientists') notion of law, shouldn't we expect that notion to reflect or be shaped by the information that we actually have access to in reasoning about laws, rather than information we do not have access to?

M: I don't deny that scientists often reason on the basis of the considerations you describe. However, it is obvious that it must be possible in principle to reconstruct what they are doing so that it fits the BSA. For example, if we take one of these assumptions - - call it p - - with modal content that you are talking about and test it by making use of another assumption with modal content q and so on, then the worry is that we either (i) are led to a regress (testing assumption q requires modally committed assumption r , testing r requires modal assumption s and so on) that never ends up being grounded just in what is non-modal or else (ii) we end up with a circle (having to assume p again).

Neither possibility is satisfactory from the point of view of justification. The only remaining alternative is that the modal assumptions on which we rely must ultimately be justified in some other way, and this can only involve some sort of balancing of simplicity and strength. Thus the inferences you describe ultimately must be reconstructable within the BS framework.

JW: I don't understand why you are so confident that it must be possible to carry out this reconstruction.

M: If it is not possible to do so, you will be relying on modal claims that are not reducible and reasoning involving such claims will be unclear, involved in a regress, and circular.

JW: First, I'm puzzled about why you phrase things in terms of whether it is *possible* to provide such a reconstruction or grounding or whatever, rather than in terms actually producing the reconstruction. If there is worry about modal claims being unclear and so on, why don't we need to actually exhibit the details of the reconstructions/reductions you describe in order to assuage these worries¹²? Or at least provide enough of the details to make it plausible that such reconstructions are always going to be possible? But you've just agreed that the scientists employing modal claims do not in fact actually ground them in the way you describe and apparently don't feel any need to produce reductions of the modal to the non-modal. Moreover, you haven't yet pointed to any epistemic or conceptual calamities that ensue from these omissions. Indeed, if there were such calamities, why wouldn't they appear in scientific practice and lead scientists themselves to be concerned to address them, which I take it that we both agree does not happen? In the absence of such defects, why isn't the obvious moral that the sort of reduction you are talking about just is not required if what one is interested in is elucidating how scientists and others provide evidence for, test, and reason about causal and nomological claims?

I can see why, given the way that metaphysicians conceive of what they are doing, there is a motivation within metaphysics for requiring that grounds/truth makers and perhaps reductions be produced, but it sounded earlier as though you were claiming something more than that—that such truth-makers need to be provided if scientific inquiry is to avoid circles and regresses, or to “make sense” or avoid fatal unclarities or be adequately grounded at some deeper level or something like that¹³. Is your view

¹² And correspondingly to recognize that some modal claim is unclear or illegitimate, don't we have to show that the required reduction cannot be produced? This would presumably require consideration in some detail of various ways that the reduction might go and showing that none of these succeed.

¹³ Might we think of Metafisico's view as something like the following: scientists don't provide reductions of the modal notions they use and typically have no opinions about (and don't seem to care about) whether it is possible to do so. Nonetheless (i) it *is* possible to do so and (ii) it is only because this is possible that the use of (unreduced) modal notions in science is legitimate. As nearly as I can see, there is nothing incoherent

perhaps that doing metaphysics but not science requires providing grounds and truth-makers? I'm happy to agree with that, as long as you agree that I am not required to do metaphysics

I'll add that, given that, as I think you have to acknowledge, the notions of "simplicity" and "best balance" (and, for that matter, "strength") that figure in the BSA analysis are not exactly pellucid, I don't see that a reduction of modal notions via the BSA, even if it could be provided, would somehow endow these notions with a "clarity" and "intellectual respectability" that they would otherwise lack, at least in any ordinary sense of those words. I think perhaps you are confusing "sufficiently clear and unambiguous to guide reasoning, at least for certain purposes and contexts" with a certain conception of "reductive".

M: This is getting tedious. As I have been trying to explain, the cogency and legitimacy of the demand for reductive truth-makers for modal claims is acknowledged by many of the greatest philosophers, past and present. It is central to metaphysics. Only an ontological philistine would resist this demand.

JW: OK. Maybe I could get a better handle on what the BSA involves if we could discuss how it bears on another issue that I'm interested in. This has to do with the status of the various generalizations that figure in the so-called special sciences. These generalizations seem to have limited domains of application and typically have exceptions even within the domains in which they are intended to apply. Nonetheless we make distinctions among these generalizations—some seem to describe causal relations and are accorded a central role in theory construction and causal explanation while others, although no more exception-ridden, are not regarded as having a different status. For example, it is uncontroversial that changes in the money supply are usually correlated with changes in the inflation rate—there is a true or approximately true generalization describing this correlation. But in the comparatively recent past there was considerable disagreement among economists about whether changes in the money supply cause changes in the inflation rate or whether instead the causality runs in the opposite direction from inflation to money or whether instead the correlation between these two factors is due to some third factor. What does the BSA tell us about the difference (if any) between the claim that money and prices are correlated and the claim that money causes prices? If it is true that money causes prices, would it be appropriate to think of this as a *law* of economics as some philosophers of social science have suggested or should we be thinking about it in some other way? To take another example, if I was an economist contemplating the Philips curve describing the inverse relationship between inflation and unemployment in the period before 1970, would the BSA analysis tell me if I should regard this generalization as an economic law? Or whether we should think of it as describing a causal relationship rather than a mere correlation?

M: The status of generalizations in the special science is an interesting question and one that is relatively unexplored within the BSA, which has focused mainly on laws in

about this position. Nonetheless, it offers (and I see) no reason to believe either (i) or (ii). Thanks to Matthew Slater for some helpful comments here.

fundamental sciences like physics. Still several suggestions have recently emerged and are the subject of ongoing discussion. It is probably most useful to begin with the version of the BSA formulated by Lewis (e.g., 1999) and others. On this version, it seems doubtful whether there are any laws in the special sciences. For one thing, recall that the BSA characterizes laws in terms of a notion of simplicity that is tied to the perfectly natural properties; the laws will be generalizations that are simple when framed in terms of such properties. Presumably the perfectly natural properties are, at least to a large extent, those that figure in fundamental physics. It is unlikely that any of the properties that figure in the generalizations of the special sciences are perfectly natural. When formulated in terms of perfectly natural properties from physics, the generalizations of the special sciences are likely to be horrendously complicated and non-simple. On this ground alone, they are disqualified as laws. To this it may be added that, often if not always, whether these special science generalizations hold seems to depend on whether various physically contingent initial and boundary conditions hold. These initial and boundary conditions are likely to be very non-simple and to contribute relatively little to strength, since in many cases they will hold only very locally—e.g., in the case of the psychology and economics perhaps only for human beings and a few other organisms. So those initial/boundary conditions are unlikely to make it into the Best System. Hence quite apart from the point about non-natural properties, the generalizations of the special sciences are unlikely to be derivable as theorems of the best system and hence are not laws.

JW: I see. I'm afraid, though, that this doesn't seem to help with the problems I described. As I said, there does seem to be distinction between those true generalizations in the special sciences that describe causal relationships (and, at least according to me¹⁴, have other features like stability or invariance under changes) and those true generalizations that do not have these features. These seem like distinctions that are tracking something in the world. That is, it looks to me as though there is a difference between, on the one hand, the way the world is if it is true that money causes prices and, on the other hand the way that it is if money does not cause prices. And there are important issues about the sort of evidence that would be relevant to establishing which if either of these claims is true – that is why the social and behavioral sciences are full of procedures for inferring and testing causal claims. Even if it is correct that there are few laws in the special sciences – a claim that I have defended myself -- if the application of the BSA to the special sciences classifies all true generalizations in the special sciences into an undifferentiated group as “non-laws” it does not seem to help with the issues that interest me. Put differently, on this version of the BSA, the law/non-law distinction does not seem to correspond to or to illuminate the causal/ merely correlational distinction that I would like to understand.

M: The version of the BSA we've been discussing may not help with the problems you describe, but does have a consequence that is far more metaphysically important: if correct, it shows that there are no laws outside of physics. This is a claim about what the world is like on a fundamental level and thus of the greatest ontological significance.

¹⁴ Cf. Woodward, 2003.

Surely you don't claim that the question of whether there are any laws outside of physics is unimportant?

JW: I don't think that asking whether the generalizations found in disciplines outside of physics are "laws" is the most useful way of formulating issues about the nature and status of those generalizations. In general, I think that you metaphysicians focus far too much on the notion of law. I do think there are facts about whether various relationships studied in the special sciences are stable in various ways, about whether these can be exploited for purposes of manipulation and control and it is these features that philosophers of science should be interested in. I don't much care about whether you call these "laws". Refusing to call these relationships "laws" doesn't make them disappear or make them unworthy of study. On the other hand merely bestowing the honorific "law" on them doesn't do much to elucidate them either.

M: Whatever. If you don't like the consequences for the special sciences of the version of BSA just sketched, here is a variant on this account which leads to results that you may find more appealing. It is called the Better Best Systems account (Cohen and Callender, 2009) and in effect it relativizes the BSA to each of the special sciences¹⁵. Take some special science of interest—e.g., economics or ecology with its proprietary predicates and its associated Humean Basis described in terms of these predicates. Consider alternative systemizations $\{S_E\}$ of this Basis. The laws of the science in question are then just those axioms (or axioms and theorems) that are common to the systemizations in $\{S_E\}$ that achieve a best balance of simplicity and strength. In fact, we can think of this framework as supplying the underlying metaphysics that is missing from your talk of invariance and relations exploitable for purposes of manipulation and control. Moreover, the respectability of the special sciences is restored, since we have shown (or it least we have shown how it might be possible, which is all that matters in metaphysics) that they contain laws. This result should please you, as someone who is interested in these areas of inquiry and seems to take them more seriously than they probably deserve.

JW: I don't mean to sound ungrateful, but I still have the feeling that the problems that interest me still haven't been addressed. If I have understood correctly, you would like to conceptualize the question of whether money M causes prices P as a claim about whether there is a law linking M and P . This law claim will be true or not depending on

¹⁵ In the interest of fairness and completeness (and to avoid leaving a misleading impression about his views) I should point out that Callender has been quite critical of some aspects of work in contemporary analytical metaphysics. See his (2011). It is also true, as he has emphasized to me in correspondence, that many advocates of the BSA present it as an anti-metaphysical view of laws or at least an approach that carries minimal metaphysical commitments with it. I agree that it is a virtue of the BSA that it avoids the kind of metaphysics that postulates non-naturalistic special entities and relationships. On the other hand, at least from my perspective, the BSA retains some important metaphysical commitments, including its aspirations to provide a reduction of the modal to the actual and its reliance on ideas about "perfectly natural" predicates. It remains more metaphysical than some of its advocates realize.

whether this generalization is an axiom or theorem in the BBS. If instead, some appropriate generalization linking P to M was such an axiom or theorem, then it would be the case that P causes M. Unlike the BS, the BBS thus at least allows for the possibility that one of these claims might turn out to be a law.

M: That's right.

JW: I still don't see how to *use* the BSA to determine which of these alternative law claims is correct, especially since the notions of simplicity and best balance on which you are relying have been left so vague and underspecified. Maybe it would help if I posed matters this way: Econometricians make use of various sorts of tests ("causality tests") and appeal to certain kinds of evidence to try to determine whether M causes P or vice-versa. What does the BSA have to say about these? It just does not seem very helpful in clarifying what the disagreement in this case is about or what sort of evidence is relevant to settling it. Does the application of the BSA to the special sciences show which of these "causality tests" are good ones for identifying causal relationships? Can we actually use it to distinguish genuine causal relationships from mere correlations in social science contexts?

M: It may well be true that the BBS can't be used to determine which of these competing causal claims is correct if that is what you want to know. But *if* it is true that e.g. money causes prices, the BBS can be used to capture or represent the metaphysical status of this generalization in a satisfactory way, rather than leaving it obscure and mysterious, ontologically speaking, which is what you have done¹⁶. More generally, I certainly wouldn't expect metaphysics to be "useful" in the sense you seem to have in mind or for metaphysical treatments of causation to cast light on issues of evidence and testing. Perhaps you might explain more clearly what it is that you are looking for. I've always thought that a sharp separation of what is merely pragmatically useful (or of merely epistemic or methodological significance, as considerations having to do with evidence and testing are) from the point of view of the parochial interests of human beings from issues about what is metaphysically true is one of the central achievements of contemporary analytic metaphysics. The Sublime Uselessness of our subject has always been one of its most appealing characteristics for me.

JW: Here is an illustration of what I have in mind in talking about "usefulness". As I mentioned earlier, the interventionist approach to causation that I favor associates causal

¹⁶ The reader may suspect that what is going on in this example is that a judgment about whether money causes prices or vice-versa is being made on some independent basis (i.e., some basis independent of explicit reliance of the BBS) and that this judgment is then being used to support the assertion that $(M \rightarrow P)$ rather than $(P \rightarrow M)$ is a theorem (involving a best balance of simplicity and strength) of the BBS, so that the latter assertion is, so to speak, an after the fact exercise in labeling or rationalization. I can only say that I sympathize with this reaction.

claims with claims about what the outcomes of hypothetical experiments would be, were an intervention to occur on the putative cause variable. I see this as having a number of implications for the methodology of constructing and testing causal claims. Let me focus on the latter— testing was not the main focus of my book, but it may help to suggest what I am driving at when I claim that an interventionist account of causation of the sort that I favor can be methodologically useful and when I worry that the sort of metaphysical foundations you are looking for are not likely to be useful for these sorts of purposes. Suppose that we are interested in using non- experimental evidence E , in conjunction with other information I , to assess whether X causes Y is true. Then the question we should ask ourselves, according to interventionists, is whether E and I provide reason to think that, if a properly controlled experiment were to be performed in which X is manipulated, Y would change. In other words, in cases of causal inference from non-experimental data, we should think of ourselves as trying to infer from that data what the results of an experiment would be without actually carrying out the experiment-- we use the hypothetical experiment to clarify what would be required to establish the causal conclusion on the basis of non-experimental information. In fact, there are formal results about this: it is possible to show that, given certain natural assumptions, we can infer that Y will change under an intervention on X from appropriate non-experimental data about correlations, given certain background assumptions. For example, we can infer to this conclusion if there is some third variable Z which is (i) known to cause X , (ii) known to cause Y if at all only through X and (iii) which is such it is uncorrelated with other causes of Y (except for those variables if any on the causal route from Z to X to Y). A variable meeting these conditions is called an *instrumental variable* by econometricians and if Z is such an instrument, and X and Y are correlated we can infer that X causes Y . Indeed, we can provide a quantitative estimate of the “treatment effect” of X on Y from information about the covariances of X , Y and Z : effect of X on $Y = Cov(Y,Z)/Cov(X,Z)$ (cf. Winship and Morgan, 1999).

The relevance of this to interventionist treatments of causation is that the underlying logic is readily understandable in interventionist terms: one can show that under conditions (i)- (iii), Z satisfies (modulo a complication¹⁷) my conditions for an intervention on X with respect to Y ; thus (according to the interventionist

¹⁷ The complication is that although it follows from conditions (i)- (iii) that conditions (I1), (I3) and (I4) in my characterization of an intervention (Woodward, 2003, p.98) are satisfied, it does not follow from (i) – (iii) that my remaining condition (I2) (that the intervention act as a switch or arrow-breaker with respect to X) is satisfied; indeed in many real-life cases instrumental variables will not be switch-like. Fortunately, Eberhardt (2007) in effect has shown that one may replace (I2) with a weaker condition (characterizing a so-called soft intervention) which does not require that interventions be switch-like but only that they supply an exogenous source of variation in the variable intervened on that is uncorrelated with the other causes of that variable. This notion of a soft intervention enables many of the same inferences as one can make in the presence of arrow-breaking interventions and also can be used to define various causal notions characterized in Woodward, 2003. Instrumental variables are very often interventions in this weaker sense.

characterization of causation) if X and Y covary under an intervention on X , then X causes Y .

There are many, many other examples of how thinking of problems of causal inference in non-experimental contexts as though one were trying to ascertain what the results of a hypothetical experiment would be without actually doing the experiment can help to elucidate what sort of evidence is required for reliable causal inference and can disambiguate and clarify competing causal claims

I'll add that I don't mean to single out the BS or BSS in suggesting that they don't seem very relevant to epistemological or methodological issues associated with causal claims and laws. Saying that the truth makers of causal claims (or laws) are facts about dispositions or relations of necessitation between universals seems equally unilluminating for purposes of understanding how can use evidence to assess causal claims, use such claims reliably in reasoning, and so on.

M: Once again, I think you misunderstand the goals that we metaphysicians have in looking for truth makers and foundations. We are not trying to address issues of non-experimental or experimental design or to evaluate statistical techniques for testing causal claims. We are looking for something much deeper -- an understanding of what is going on at the level of fundamental metaphysics.

JW: Fair enough—I get that you are not trying to do methodology and also that you think that metaphysics of science should be done in a way that abstracts from all merely methodological or practical concerns. What I do *not* get is why you also insist that I have to provide an account of the grounds or metaphysical foundations for causal claims before I can address the epistemological and methodological issues that interest me. On the contrary, given our discussion so far, it looks to me as though even if we had the right story about the grounds or truth makers for causal claims, this would not help with the issues that interest me.

M: Metaphysics does indeed bear an asymmetric relationship to the other topics you mention, but this should not be surprising in view of its foundational role. Briefly, epistemology and methodology are not at all relevant to metaphysics but metaphysics is highly relevant to them, insofar as they have any philosophical interest. However, the details of how this works does not matter for present purposes and, in any case, is not the concern of the metaphysician, as one can see by examining virtually any contemporary discussion of the metaphysics of causation. But rather than continuing in this vein, let me try another tack. I take it from what you said above about interventions that a necessary condition for I to count as an intervention on X with respect to Y is that intervention I cause X .

JW: That's right.

M: But then your account, which characterizes causation in terms of the response of Y to interventions on X is blatantly circular.

JW: I acknowledged in *MTH* that my account is non-reductive, although I also pointed out that the causal information that figures in the characterization offered of what it is for X to cause Y has to do with other relationships besides the $X \rightarrow Y$ relationship. In other words, the characterization suggests how you can use information about other causal relationships to assess whether X causes Y , but does not require that you have to already know whether X causes Y to use the characterization. Thus the circularity involved in my characterization, such as it is, is not vicious, at least for epistemic purposes—indeed, at the level of particular causal claims, there is no circle at all, although there is also no reduction of those claims to claims that are non-causal¹⁸. Incidentally, the example involving instrumental variables above provides a concrete illustration of this point—use of this method (when understood as above) requires prior causal knowledge but it can be used to provide evidence for different causal relationships. So a non-reductive account can be quite useful for epistemological or methodological purposes. Again my conjecture, based on our conversation so far, is that accounts of laws and causal relationships that are reductive in aspiration are particularly unlikely to be helpful in clarifying the methodological and epistemological issues that interest me, both because the notions that figure in reductive accounts (simplicity and so on) are so unclear and because in the interests of making the reduction work philosophers seem to end up misrepresenting what it is they are trying to reduce, as with the idea that good scientific method involves a trade-off of simplicity and strength. I think that a similar point holds for the non-reductive accounts of truth makers that you have mentioned.

M: I'm afraid that you completely lack an ear for metaphysics. You are one of the most ontologically shallow people I have ever met. (Aside: It is not vouchsafed to just anyone to discern the lineaments of Being.)

JW: I fully agree. Let me buy you a drink.

References

Armstrong, D. (1983) *What Is a Law of Nature?* Cambridge: Cambridge University Press.

Batterman, B. (1998) "Why Equilibrium Statistical Mechanics Works: Universality and the Renormalization Group," *Philosophy of Science*, 65, 183-208.

¹⁸ The topic of circular definitions deserves much more attention than I can give it here. As Anil Gupta has long argued (e.g. 2014), there are any number of "definitions" in mathematics and science that are formally "circular" but nonetheless provide useful information and illumination. I owe to Mark Wilson the following example: i = the unique number x such that $x^2 = -1$ & $x = i$ (personal communication); this is "circular", yet captures the core behavior of the complex numbers. In general, even if "circular", a definition can function so as to exclude some cases from falling under the defined term and include others

Bird, A. (2007) *Nature's Metaphysics: Laws and Properties*. Oxford: Oxford University Press.

Callender, C. (2004) "Measures, Explanation and the Past: Should 'Special' Initial Conditions Be Explained?" *British Journal for the Philosophy of Science* 55, 195-217.

Callender, C. (2011) "Philosophy of Science and Metaphysics" in S. French and J. Saatsi, (eds.) *The Continuum Companion to Philosophy of Science*. London: Continuum International Publishing Group

Chalmers, D., Manley, D. and Wasserman, R. (2009) *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press.

Cohen, J. and Callender, C. (2009) "A Better Best System Account of Lawhood" *Philosophical Studies* 145(1):1-34.

Correia, F. and Schnieder, B. (2012) *Metaphysical Grounding: Understanding the Structure of Reality*. New York: Cambridge University Press.

Dretske, F. (1977) "Laws of Nature" *Philosophy of Science* 44:248-68.

Eberhardt, F. (2007). *Causation and Intervention* (Ph.D. Thesis). Carnegie Mellon University. Available at philosophy.artsci.wustl.edu/people/Frederick_Eberhardt. Accessed 8/12/11/

Goldenfeld, N. and Kadanoff, L. (1999) "Simple Lessons from Complexity" *Science* 284: 87-9.

Gupta, A. "Definitions" *Stanford Encyclopedia of Philosophy* <http://plato.stanford.edu/entries/definitions/> Accessed, October 26, 2014.

Holland, P. (1986) "Statistics and Causal Inference" *Journal of the American Statistical Association* 81:945-60.

Kendler, K. (2005) "'A Gene for..': The Nature of Gene Action in Psychiatric Disorders". *American Journal of Psychiatry* 162: 1243-52.

King, G., Keohane, R. and Verba, S. (1994) *Designing Social Inquiry*. Princeton: Princeton University Press.

Lewis, D. (1999) *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.

Loewer, B. (2009) "Why Is There Anything but Physics?" *Synthese*. 170: 217-233:

Manley, D. (2009) “Introduction: A Guided Tour of Metametaphysics” in Chalmers, D., Manley, D. and Wasserman, R. (2009)

Morgan, S. and Winship, C. (2007) *Counterfactuals and causal inference: Methods and principles for social research*, Cambridge: Cambridge University Press.

Mumford, S. (2004) *Laws in Nature*. London: Routledge

Pearl, J. (2009) *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.

Reiss, J. (2005) “Causal Instrumental Variables and Interventions” *Philosophy of Science*, 72: 964–976.

Rubin, D. (1974) “Estimating Causal Effects of Treatments in Randomized and Non-Randomized Studies” *Journal of Educational Psychology* 66:688-701.

Schaffer, J. (2009) “On What Grounds What” in Chalmers, D., Manley, D. and Wasserman, R. (2009)

Seider, T. (2009) “Ontological Realism” in Chalmers, D., Manley, D. and Wasserman, R. (2009)

Strevens, M. (2007) “Review of Woodward, Making Things Happen” *Philosophy and Phenomenological Research* 74 (1):233–249.

Tooley, M. (1987) *Causation: A Realist Approach*. Oxford: Oxford University Press.

van Inwagen, P. (2007) “Metaphysics” *Stanford Encyclopedia of Philosophy*. Accessed, July 31, 2011.

Woodward, J. (2003) *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.

Woodward, J. (2014) “Simplicity in the Best Systems Account of Laws of Nature” *British Journal for the Philosophy of Science* 65:91-123