Comp. by: PG3754 Stage : Proof ChapterID: 0001565204 Date:12/6/12
Time:19:58:24 Filepath:d:/womat-filecopy/0001565204.3D208

10

Causation, Free Will, and Naturalism[1]

Jenann Ismael

> Concepts which have proved useful for ordering things easily assume so great an authority over us, that we forget their terrestrial origin and accept them as unalterable facts. . . . It is therefore not just an idle game to exercise our ability to analyze familiar concepts, and to demonstrate the conditions on which their justification and usefulness depend, and the way in which these developed . . . in this way they are deprived of their excessive authority.
>
> (Einstein quoted in Born, 1971: 159)

The problem of free will arises in a distinctive form in the Newtonian context with the provision of strict deterministic laws of nature that allow our actions to be derived from conditions in place at the beginning of the universe.[2] This is understood as meaning that no matter how much it feels like our actions are flowing freely from our choices, they are in fact causally

[1] I owe a great debt to Michael Mckenna and Keith Lehrer for very helpful comments, and to Dave Schmidtz for the opportunity to present this material in his seminar. Also, a very warm thanks to James Ladyman and Don Ross for including me in the workshop that gave rise to this volume, and the warm hospitality of the philosophy department at the University of Alabama, Birmingham. And, like all of my work from 2005–2010, to the Australian research council for support.

[2] This is just one horn of what is known as the Dilemma of Determinism. The problem doesn't disappear if the laws incorporate the sort of indeterminism characteristic of quantum phenomena. The reason is that we want our wills to be the originators of action, not random, spontaneous acts of nature.

determined by facts that were in place before we ever came on the scene. In recent years, defenders of free will in the philosophical literature have focused on the notion of personal autonomy, showing that in paradigmatic cases of willful action there are psychological structures in place that make it right for me to say that action flows from me in a morally relevant sense. The focus turned to psychological requirements for freedom and removing the kinds of impediments we all recognize as undermining the will to act: addiction, delusion, akrasia, misinformation, and coercion. These are the sort of things Denis Overbye is jokingly alluding to when he writes 'I was a free man until they brought the dessert menu around'.[3] These kinds of constraints pose real and substantial threats in everyday life. Freedom from them is not an all or nothing matter and it is not easy. There is a lot of very useful discussion in the moral psychology literature about what free will, understood in these terms, amounts to and how to attain it.

But to someone gripped by the worry that our actions are causally determined by the initial conditions of the universe, talk about morally and psychologically relevant senses of freedom can seem beside the point. We can talk until we are blue in the face about freedom from psychological constraints on action, but if our actions are *causally* necessitated by the earliest state of the universe, we are no more free to act otherwise than a leaf is free to blow against the wind, or a river is free to flow up the mountain rather than down. Of course, we don't typically know the causal determinants of our actions, but that doesn't mean that we are causally capable of acting otherwise than we do.

This is a worry that need not deny that you may act as you choose, but it points out that your choices themselves have causal antecedents. The mere existence of causal antecedents to your choices means that even if you satisfy all of the requirements for personal autonomy, you are causally compelled to act as you do. Here are some expressions of this worry:

> Causal determinism entails that all human choices and actions are ultimately compelled by . . . a complex chain (or web) of causal antecedents. (Ferraiolo 2004: 67)
>
> If determinism is true, . . . all those inner states which cause my body to behave in what ever ways it behaves must arise from circumstances that existed before I was

[23] New York Times [online] <http://www.nytimes.com/2007/01/02/science/02free.html> accessed 17 May 2012.

born; for the chain of causes and effects is infinite, and none could have been the least different, given those that preceded. (Taylor, 1963: 46)[4]

I want to confront this worry head on, starting with the pre-theoretic view of causation, showing how what we've learned from science about the deep logic of causal talk undermines the conception of cause that makes this seem like a worry.

The plan for the chapter is as follows: I will start with the folk notion of cause, lead the reader through recent developments in the scientific understanding of causal concepts, showing how those developments undermine the threat from causal antecedents, and end with a happy complicity of a scientific vision of the world that in no way undermines the happy confidence in one's own powers to bring things about. Then I'll make some methodological comments, using the discussion here as a model for a kind of naturalistic metaphysics that takes its lead from science, letting its concepts be shaped and transformed by developments in science.

The problem of free will is a notorious Gordian knot, a mess of tangled threads that need to be teased apart and addressed in turn. Each one of these threads holds important lessons for our understanding of ourselves and our place in nature. The worry about causal necessitation is only one of these threads, but one that holds some important lessons about the nature of natural necessity.

The evolution of causal notions

The concept of cause has always played a central role in people's thinking about the natural world, but it has undergone a quiet transformation in science that not all philosophers are aware of. Causal thinking arises spontaneously in the child at about the same time that it begins to recognize that it has willful control over its limbs. That process of discovering the robust connections that allow us to act effectively continues into adult life where knowing the causal effects of our actions in the short and long term is essential to effective agency. The most rudimentary ideas about cause and effect have to do with the fact that by exerting various forces and doing work on other material systems, one could make them behave as one

[4] Taylor is not endorsing this view, but formulating what he calls the standard argument against free will.

wants. In this incarnation, causes are closely connected with mechanical concepts like force and work. Science can be seen as developing out of systematization and abstraction of causal thinking, the search for an understanding of the causal relations among events. At first, the notion of cause in science retained its close connection with mechanical ideas. A cause was something that brought about its effect by a transfer of force. But in Newton's physics, this connection is loosened. Physical laws take the form of differential equations that give the rate of change of a quantity over time. Force drops out of the picture. The idea of compulsion or even a *direction* of influence is lost. As David Bohm put it:

> It is a curiously ironical development of history that, at the moment causal laws obtained an exact expression in the form of Newton's equations of motion, the idea of forces as causes of events became unnecessary and almost meaningless. The latter idea lost so much of its significance because both the past and the future of the entire system are determined completely by the equations of motion of all the particles, coupled with their positions and velocities at any one instant of time. (Bohm 1989: 151)

Russell thought these sorts of relations embodied in dynamical laws were so different from causal relations as traditionally conceived, that it is misleading to think of them in causal terms at all. He gave two main reasons for rejecting a causal interpretation of the dynamical laws. The first reason was what Bohm was remarking on, viz. that causal relations incorporate a temporal asymmetry that dynamical laws do not. They are what are sometimes called regularities of motion that relate the state of the world at one time to its state at others, but there is no direction of determination. There is no suggestion that one state *produces* another, brings it about, or compels its occurrence. It is true that fixing earlier states fixes later states. But it is also true that fixing later states fixes earlier ones. The dynamical relationship is entirely symmetric and incorporates no direction of determination.

The second reason that Russell gave for rejecting a causal interpretation of the dynamical laws was that causal relations hold between relatively localized events at a temporal distance from one another, like the striking of a match and the appearance of a flame, or the turning of a car key and the starting of an engine. The dynamical laws, by contrast, relate only states of the world as a whole (in relativistic physics, spatial hypersurfaces). This is connected to a third difference between dynamical laws and causal generalizations that others have noted. Causal generalizations are imprecise and defeasible. We all know that to get matches to light and cars to start, all kinds of subsidiary conditions—not usually mentioned and not always explicitly known—have to be in place. Things go wrong, and we don't always know what. Dynamical laws, by contrast, don't allow exceptions. They hold with perfect precision at all times and places. Unfortunately, because they relate global time slices of the world, they are of little use for the purposes of guiding action or prediction at the local scale. Although they are less precise and always defeasible, causal generalizations are much more useful for beings with limited information and local input to the world.

Russell's view in the 1918 paper was that causation is a folk notion that has no place in mature science. As he (famously) says:

> The law of causality, I believe, like much that passes muster among philosophers, is a relic of a bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no harm. (Russell 1918: 180)

There are few people who agree with Russell's conclusion, but his remark set off a line of questioning about the status of causal relations. What we've learned, particularly in the last twenty years, about the logic and content of causal claims casts the problem of causal determinants in a rather different light. The modern discussion of causation in the philosophy of science really began with Cartwright's deeply influential and telling criticism of Russell's paper (Cartwright 1979). Philosophers often cite Gettier as a rare example of someone who genuinely refuted a philosophical view, but Cartwright's argument against Russell came quite close. She pointed out that dynamical laws cannot play the role of causal relations in science because specifically causal information is needed to distinguish effective from ineffective strategies for bringing about ends. So, for example, it might be true as a matter of physical law (because smoking causes bad breath), that there is a strong positive correlation between bad breath and cancer. But it is not true that bad breath *causes* cancer and hence it is not true that treating bad breath is an effective strategy for preventing cancer. And that difference—the difference between being correlated with cancer and being a way of *bringing it about*—is not one that can be drawn by looking only at dynamical laws. If one wants to avoid getting cancer, one has to know not simply what cancer is correlated with, but what *causes* it, that is, what brings it about.

There was a lot of handwringing, wondering what causal information adds to dynamical laws. Philosophers entertained probabilistic and counterfactual analyses, and there were a lot of unresolved questions about the metaphysics of causal relations. In the last fifteen years, there has been a quiet revolution in how we model, understand, and learn about the causal structure of the world. The revolution started in philosophy and computer science, but spread to psychology and statistics where theories of causal learning and statistical inference have made huge steps. For the first time, we have available a comprehensive formal language in which to represent complex causal systems and which can be used to define normative solutions to causal inference and judgment problems (Pearl 2000).

New insights are emerging into the role of causality, causal models, and intervention in the basic human cognitive functions: decision-making, reasoning, judgment, categorization, inductive inference, language, and learning. Facts about how people informally think, talk, learn, and explain things in causal terms, are formalized and systematized in scientific contexts, which makes an art of causal modeling. There are complementary formalisms (graphs and counterfactual analyses, Bayes nets), but the interventionist account has emerged as a clear forerunner. The formal apparatus gives us ways of rendering the deep causal structure of situations, refines intuitions and gives us positive criteria for making assessments in hard cases (Pearl 2000).

The central idea of the interventionist account is that causal structure encodes information about the results of hypothetical interventions. 'Intervention' is a term of art that refers to interaction that effectively randomizes the value of a variable, so that it can be treated probabilistically as a free variable. Although interventions are supposed to be defined formally and independently of human action, human actions turn out to be an important class of interventions,[5] and so the role that causal information plays in guiding strategic action (noticed by Cartwright in her objections to Russell) fall very nicely out of the interventionist account of the content of causal claims. We need causal information to decide how to act, because we need to know how our actions will affect the future, quite independently of how they might themselves be affected by their own pasts.

The interventionist account came out of independent work by Glymour's group at Carnegie Mellon and Judea Pearl at UCLA. Pearl's work culminated in his beautiful book *Causality* (2000) and became known to more philosophers through James Woodward's *Making Things Happen* (2003). It

[5] In many interactions, volition-governed action effectively randomizes the effects of external variables, breaking correlations between a variable and earlier parameters.

provides a rich formal framework for representing causal relations in science and makes it easy to express in logical terms what causal information adds to the information embodied in dynamical laws.[6] Dynamical laws tell us how the state of the universe at one time is related to its state at another and they entail a complex web of interdependencies among the values of physical parameters. What causal information adds is information about what would happen if a given parameter were separated out of this web, severing connections with antecedent variables, and allowed to vary freely. The term 'intervention' is introduced to describe the process of separating of a parameter from antecedent variables and allowing it to vary freely, hence the name and the slogan of the interventionist account: causal information tells us what *would* happen under hypothetical interventions on a chosen parameter A.[7]

The interventionist account captures quite nicely the intuitive idea that causal information separates the *effects* of a parameter from information it carries in virtue of its own connections to causes in the past. Knowing the causal effects of A is knowing how the values of downstream variables are affected by free variation in A. Even though there may be a strong law-like correlation between having breath that smells like cigarettes and developing cancer, bad breath is not a *cause* of cancer if the cancer rate would not be altered by giving people mints to improve their breath. Smoking, by contrast, is a cause of cancer if abolishing smoking would lower cancer rates. Intuitively, causal structure separates what a parameter *does*—namely, the effects that it brings about—from the information it happens to carry about the future in virtue of common causes in the past. Since that is a difference that only shows up when its own links to other variables are severed, and since no variable is ever actually severed from its causes, this extra content can be captured only in counterfactual—or, as interventionists like to say, hypothetical—terms. Think of a newscast in which Obama,

[6] By 'dynamical laws' here, I will always mean fundamental laws of temporal evolution, the kinds of things expressed by Newton's equations and defined, in the first instance, for the universe as a whole.

[7] It turns out to be subtle to characterize interventions explicitly and there are disagreements among interventionists about the right formal definition. But there is agreement that intervention is an indisputably causal notion. Interventions can be characterized as a special class of causal processes, but not in non-causal terms. So the interventionist account aims for elucidation of the truth-conditional content of causal claims (the inferential relations among them and the conditions and consequences of accepting them), but doesn't reduce causal facts to non-causal ones.

announcing new measures in Afghanistan, follows a weatherman forecasting rain. Firing the weatherman won't ward off rain, but firing Obama would ward off the announced measures in Afghanistan, but that difference—the difference between Obama's connection to the events he announces, and the weatherman's connections to the weather—is a specifically causal one, one that cannot be made out in terms of correlations between reports and the events reported.

One of the most interesting consequences of this account from a philosophical point of view is that causal information turns out to be *modally richer* than the information contained in fundamental dynamical laws by which I mean that it has modal implications that strictly outrun the modal implications of the laws. The laws tell us how the state of the world varies over time and hence how history as a whole would vary with different choices of initial state, but they don't have any direct implications about what would happen under free variation of local parameters that occur later in history. And this is the information we need to make causal judgments. For, once the initial conditions are given, the counterfactuals we need to assess to capture local causal structure—so-called 'intervention counterfactuals'—are counter*legals*.

It's important that it is understood that there is no logical *incompatibility* between the laws and intervention counterfactuals. The fundamental laws simply *leave out* the information contained in the intervention counterfactuals. They are *silent* about the hypothetical scenarios described in the antecedents of intervention counterfactuals. The intervention counterfactuals contain modal information that *goes beyond* the information contained in global laws; they describe hypothetical cases that the laws do not cover. This is a consequence of the formal definition. It means that for (almost) any set of global dynamical equations involving two or more variables, there will be multiple conflicting accounts of the causal structure of the system that satisfy the equations. These models will preserve the relations among the values of parameters entailed by the laws, but disagree over the results of hypothetical interventions.

A concrete example

To get a more intuitive understanding of how we should think of causal structure, let's get a concrete example in front of us. Consider a fairly familiar kind of mechanical system, say, the engine of a car. If the engine were its own little universe—that is, if it were closed with respect to external influenc—we might be able to come up with global laws that would allow us to calculate its state at one time from its state at any other. If we were just interested in predicting its behavior, this would be all we would ever need. There would be nothing unaccounted for in the model that could make a difference to its evolution, no unexpected contingencies that could divert it from its course, no action on it from the outside and no question of acting on it ourselves. It's because the engine is *not* its own little universe that there are not likely to be such laws. It is embedded in a larger universe and subject to influence from outside. There are break-downs, unexpected accidents, contingencies that can never be entirely accounted for. Any global rules that describe its behavior are defeasible regularities that come with ill-defined *ceteris paribus* clauses; claims about how it normally behaves, if all goes well.

And it is because we are not simply interested in prediction that even if there *were* such laws, we would need something more. The fact that we interact with the engine means that we are not just interested in knowing how it *does* evolve; we are interested in knowing how acting on it in various ways *would* affect its evolution. This information precedes our knowledge of how it *will* behave because it *guides* our interventions into its behavior. For these purposes, global laws that tell us how the system's state changes if left to its own devices do us little good. We need a working knowledge of the engine. Working knowledge, as we will see, is causal knowledge, and we can get a good understanding of the epistemic and practical significance of causal information by looking at what it takes to get a working knowledge of a system. Suppose Joe is an aspiring mechanic and he wants to know how engines work. It wouldn't be enough to give him a rule that allows him to calculate the state at one time from its state at another. We would tell him what the important components of the engine are and give him rules that told how they work separately and in conjunction to produce the overall pattern of evolution. Very likely, we would start with a diagram of the moving parts, which identified the valves, camshaft, piston, crankshaft, and so on. Something like this:

And then we would give him separate diagrams for each of the parts that say how their output varies with input. We might start with a model of camshaft that tells him that the camshaft transforms circular motion into a

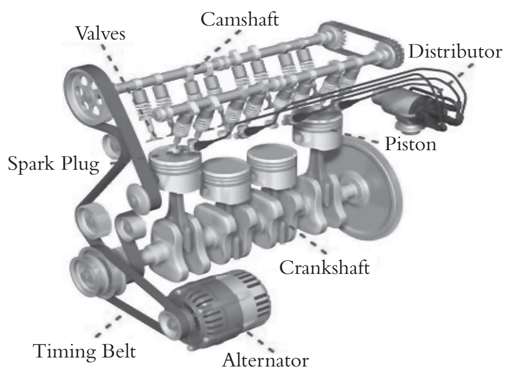

Figure 10.1 The parts of the engine in configuration.

downward force applied to the valves, and how the speed of the motion affects the strength of the force.[8]

And then we might give him a model that shows how the valves work,[9] telling him that each valve utilizes a spring, which will return them to their original position (closed) after the force is removed.[10]

He would also need a model of the piston telling him that the piston moves up and down, opening and closing the intake valve so that at the beginning of the intake stroke, the valve opens drawing the fuel-air mixture into the cylinder and when the piston reaches the bottom of the intake stroke, the valve closes, trapping the mixture in the cylinder.[11]

Once he is acquainted with how the parts behave by themselves when their input is allowed to vary freely, they are reassembled in a model that shows how the fixed connections among them constrain their relative motions, that is, how the output of one constrains the input to another. Once he has all of this, he will have what we think of as a working knowledge of the engine in the sense that he will know not only how it evolves as a whole, but also how it

[8] Chongqing CAIEC Machinery & Electronics [online] <http://www.caiecq.com/suzuki-crankshaft-for-sale-p-357.html?cPath=114_118> accessed 17 May 2012.

[9] 2CarPros [online] <http://www.2carpros.com/articles/how-camshaft-variable-valve-timing-works,> accessed 17 May 2012.

[10] 2 CarPros [online] <http://www.2carpros.com/articles/how-camshaft-variable-valve-timing-works> accessed 22 May 2012.

[11] Your Dictionary [online] <http://images.yourdictionary.com/piston> accessed 17 May 2012.

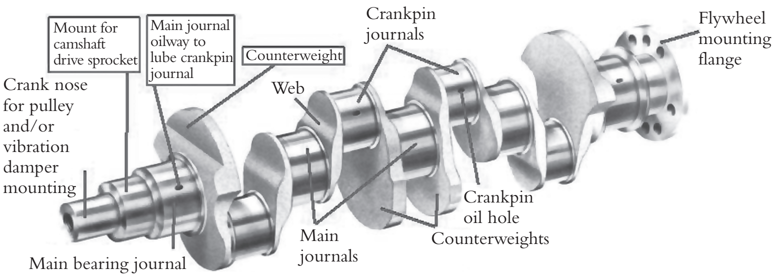

Figure 10.2 The camshaft.

decomposes into parts, the rules governing the behavior of the parts, and the mutual constraints imposed by their arrangement.

There is a reason for this procedure. There is a reason, that is to say, that we need the kinds of working knowledge embodied in the sub-models of the various parts of the system in addition to global laws of evolution.[12] The global laws contain information about how the engine evolves as a unit. The sub-models tell us what would happen if the parts were separated out and their input allowed to vary without constraint. Sub-models tell us what would happen under free variation of parameters whose values are constrained in configuration. The interventionist slogan put this point by saying that the information that is captured in these sub-models tells us what would happen if internal parameters were separated from their pasts and allowed to vary freely. I prefer to put it a little differently by saying we can't in general extract the theory of local subsystems of a larger system from dynamical laws that pertain to the whole, though the point is essentially the same. There are lots of ways of putting together subsystems to achieve the same dynamics for the overall state. Embedded in a larger machine, the input to each of the parts is constrained by the parts of the engine, so that each now moves within a restricted range and information is lost. That is why in general causal information about local subsystems goes missing when we embed them in a more encompassing model. Variables that were allowed to vary freely in the original model are constrained by the values of variables in the embedding model and so we just lose information about what *would* happen *if* they were allowed to vary freely. And again, the more encompassing model isn't *incompatible*

[12] 'Global' here means 'applying to the engine as a whole'.

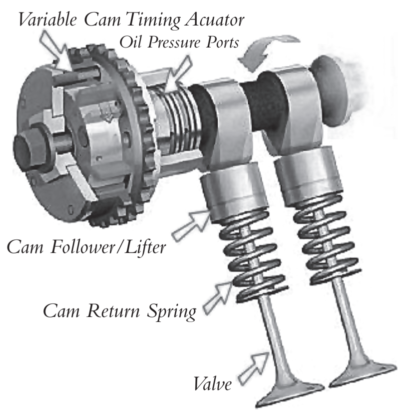

Figure 10.3 The valves.

with the less encompassing one. It is just silent about a class of hypotheticals that the second one tells us about. If we just look at the way that a system evolves as a whole, we lose information about the compositional substructure that is important to guiding our interaction with it.

The practical importance of causal information is that different ways of piecing together subsystems will have different implications for the results of local interventions. We care about causal information because our interaction with the world takes the form of changes in the values of local variables and acting strategically demands knowledge of how changing the values of local variables affects things downstream. Working knowledge matters when we are not just predicting how a system will

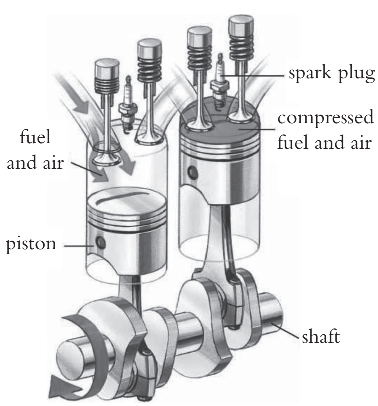

Figure 10.4 The piston.

evolve, but concerned with the question of how setting the values of internal parameters will affect its evolution. For purposes of predicting whether someone who smokes will get cancer, it doesn't mater whether smoking causes cancer or is merely a good predictor, a sign of something else that causes it. For purposes of predicting whether an engine will break down, it doesn't matter whether dirty oil causes, or is merely a sign of, impending engine breakdown. The causal information matters if you are trying to decide whether to quit smoking or whether to clean the engine pipes. It's no use doing either if they are not causally related (respectively) to cancer and breakdown.

This working knowledge of the parts of the world is essential for embedded agents interacting with open subsystems of the world. Philosophers tend to be uninterested in these partial views of bits of the world. They make a lunge for totality,[13] to get a fully encompassing view of the world as a whole as it appears *sub specie aeternitatus*. Most day-to-day science, however, is not concerned with the world as a unit, but is focused on local subsystems of the world, investigating the causal substructure of systems at different levels of description. Science raises and answers questions about what would happen under hypothetical variation of any chosen variable, holding fixed any specified class of internal connections. It was a remarkable discovery when Newton found dynamical laws expressible as differential equations that make the state of the universe as a whole at one time in principle predictable from complete knowledge of its state at any given time. But those laws are of little use to us, either for prediction or for intervention. We don't encounter the universe as a unit. Or rather, we don't encounter it at all. We only encounter open subsystems of the universe and we act on them in ways that require working knowledge of their parts.

The structure of causal sub-models

Any real subsystem of the world can be represented in countless ways: on its own, as part of an indefinite number of larger systems, at different levels of description, holding fixed different elements of internal structure, assuming different external scaffolding. There are different ways of carving things up,

[13] Milton Munitz uses this phrase to describe what he saw as Einstein's move in arriving at the theory of relativity.

and different places to draw a line between fixed and variable structure. We can model the car engine *as an engine*, which is to say, as a collection of macroscopic components bound together into a configuration that restricts relative movement. But we could equally model it as a mass of metal and rubber, holding fixed nothing about how it is arranged, or a collection of microscopic particles not even bound into metallic and rubbery configurations. And we can treat it as we did above, by itself or as part of a larger system. So how does a scientist proceed when he sets out to model a system? He creates a kind of virtual separation of a system from its environment; he puts a kind of frame around it. As Pearl says:

> The scientist carves a piece from the universe and proclaims that piece*: in* namely, the *focus* of investigation. The rest of the universe is then considered *out* or *background*, and is summarized by what we call *boundary conditions*. (2000: 15)

The variables that are included in the model are referred to as endogenous variables. The variables whose values are not included in the model are referred to as exogenous variables. He then specifies a range of values that the exogenous variables can take. This range of variability introduces a degree of freedom in the model.

He then specifies a class of invariants, features of the system that he will hold fixed for purposes of assessing the effects of variation in the values of exogenous variables. These choices affect his conclusions about the systems. When we ask how the output of the piston, or the engine as a whole varies with differences in input, it matters very much which elements of internal structure and external scaffolding we hold fixed. By scaffolding, I mean features of the external environment that are not explicitly mentioned in the model but that are crucial to supporting the fixed connections inside the system. So, for example, if we are developing a model of an engine or a building, we make some tacit assumptions about gravity, temperature, and speed. We assume the temperature is not close to absolute zero, that we are not travelling close to the speed of light, and so on. These are accounted for tacitly in the specification of the invariants. If we were travelling close to the speed of light, some of the internal regularities we want to hold fixed would break down. We don't have a well-defined model of a system until we've made choices about the exogenous and endogenous variables and the invariants. Sometimes these choices are tacit, but what we say about the system will depend on them. In modeling the engine above, for example, the conclusions we draw will depend on whether we hold fixed the internal integrity of its parts and their relative positions, or whether we want the model to cover situations in which those connections are broken.

There are different ways of specifying invariants. We can specify them directly and explicitly by saying 'hold this fixed'. In constructing a model of the engine above, if we wanted to know how the input varies with the output, we might simply stipulate holding fixed the internal configuration of parts and ambient conditions within normal range close to the surface of the earth. But we can also specify them indirectly by specifying the range of applicability of the model, and sometimes that indirect specification is also inexplicit. The normal order of discovery runs in this direction; we choose our endogenous and exogenous variables, make some assumptions—tacit or otherwise—about the range of applicability, and then test for the invariants. We just start studying a system across a range of contexts and in retrospect we find that the model doesn't apply in contexts we may only know how to characterize in retrospect.[14] But if we have a fully formulated theory, the theory will tell us what varies with changes in the values of exogenous variables across a specified range of contexts.

Once we've specified the exogenous variables and invariants (or, equivalently, exogenous variables and a range of application), we have a *causal sub-model*. These two things together will induce a separation of the properties of the system into fixed and variable structure. The fixed structure is the internal connections that remain in place across the range of applicability. Every real system will support numerous causal sub-models, each with its own range of applicability and built-in division between fixed and variable structure. And each of these causal sub-models will, in its own way, reveal something about the causal structure of the system.

The important point for our purposes is that sub-models that draw the line between fixed and variable structure in different places are not incompatible, but complementary. What is treated as exogenous in one model is treated as endogenous in another. Aspects of the system that are held fixed in one model may be allowed to vary in another. Each of these sub-models will reveal different aspects of causal structure. Causal sub-models focus attention on interesting quasi-isolated dynamical units that retain enough

[14] Before the discovery of X-rays, we wouldn't have known how to characterize contexts of high radiation. Before electromagnetic theory, we wouldn't have known how to characterize regions of electric and magnetic fields.

internal integrity under action from the outside (at least relative to a range of common contexts) to support robust, counterfactual supporting generalizations that can serve as avenues for strategic intervention. The pathways highlighted in these models are defeasible by interaction with other causes, usually supported by unrepresented features of context and often recoverable from lower-level dynamics only by simulation.[15] They play an indispensible role mediating local interaction with the world.

We have an everyday practical need for models that assess the effects of our actions on particular localized subsystems of the world—engines, toasters, computers, and cars—and, more generally, the temporally downstream effects on history of variation in parameters that correspond to decisions, in conditions that leave fixed all of the local structure that isn't directly affected by decision. Those are the models we use in deciding how to act, because we want to know how decisions will affect the history of the world. In constructing those models, we treat our own actions as free variables. And we don't typically just hold fixed the physical laws, we hold fixed all kinds of facts about our environments, all of the known and unknown infrastructure that supports reliable connections between localizable events in our surroundings and ourselves. But we also have uses for models that assess the effects of differences in family, culture, gender, or early experiences on decisions, models that assess the effects of variation in temperature on signaling in eukaryotic cells, models that assess the effect of variation in inflation rates on employment, ozone on the environment, or sugar on the behavior of naked mole rats.

Relations among causal sub-models.

There are three dimensions along which causal sub-models that represent the same system can differ from one another.[16] First, they can have different *scope*. We saw examples of this already in the relations between the sub-models of engine parts and the wider scope model of the engine in which they were embedded. Second, they can have different *invariants* (or ranges of applicability). In modeling the engine, we tacitly held fixed the internal configuration of moving parts, but we could just as easily have

[15] This is because the robust pathways are emergent structures stabilized in the feedback loop of acting on a system and observing the results of action.

[16] This may not be exhaustive.

included in our range of applicability situations in which the engine came apart, or the parts were laid on the ground apart from one another. The more we hold fixed internally, the less freedom there will be for the system as a whole, that is, the less variation there will be in its overall state. Finally, they can be at different *levels of detail* (we could describe the engine at the macroscopic or microscopic). A system that has only five components at one level of detail may have 10^5 at another. Think of what happens when we look at the engine at the level of microscopic detail.

The causal substructure of a system is only partially captured in any one of the myriad number of different ways of modeling a system, each with its own choice of exogenous variables and fixed set of internal connections. Fully specified, it should furnish the basis for judgments about what would happen under free variation of any chosen parameter, holding fixed any specified class of invariants. When we assess the effects of variation in a parameter, we are usually looking forward in time, but we can also raise questions about the effects of variation in current or future states on the past. And, there is no intrinsic restriction to localized interventions. We are often interested in assessing localized interventions, but we can certainly raise questions about the effects of variation in complex and distributed events like climate or water mass on global variables. And when we assess the effects of one variable on another, we are almost always tacitly holding fixed aspects of the internal structure of the system, and features of the environment not explicitly mentioned and in some cases not known.

We make all sorts of choices when we construct a sub-model. These choices are governed by the purposes to which the model will be put, and the modal conclusions that we draw about the system are *conditioned* on these choices. What would happen if the value of gravity were allowed to vary within a given range? What would happen if the size of the earth were increased or decreased along a given scale? Or if the wind reached speeds of 200 mph in Los Angeles? Answers to each of these questions depends on which elements of external scaffolding and internal structure we hold fixed and the range of circumstances we are talking about. Those have to be specified before we have a well-defined question.

In *actuality* every variable has a value. Modality is introduced when the value of a variable is allowed to vary over a certain range. The variability in every case is *purely hypothetical*. By letting exogenous variables take a range of values, a model introduces a dimension of hypothetical variation into the state of the world and by specifying a class of invariants, it induces a division between fixed and variable structure. We can talk about hypothetical variation in the values of variables at any point in the world's history, early or late. There is no fact of the matter about what hypothetically happens. There are only facts about what happens under certain hypothetical suppositions.

Which features of our models depend on the choice of exogenous variables and invariants, that is, which facts are *relative* to the choices that define a sub-model? The asymmetry between the exogenous and endogenous variables introduces a *direction* of influence. We tend to hold the past fixed and let the future vary, because questions about how later states vary with differences in early ones have a special importance for purposes of guiding action, but formally, there is no difficulty in asking about the effects of variation in future states on the past. So the asymmetry between cause and effect, the order of determination, is an artifact of the choice of exogenous and endogenous variables. As Pearl says:

> This choice of [endogenous and exogenous variables] creates asymmetry in the way we look at things, and it is this asymmetry that permits us to talk about 'outside intervention', hence, causality and cause-effect directionality. (2000: 350)[17]

And again,

> The lesson is that it is the way we carve up the universe that determines the directionality we associate with cause and effect. Such carving is tacitly assumed in every scientific investigation. (2000: 350)

And it is not just the *direction* of influence. The *existence* and *strength* of influence depends on what we hold fixed and what we allow to vary. Whether A and B are connected at all and the strength of that connection and how depends on what we hold fixed. Even the most robust local connection like the connection between smoking and cancer, or even the ambient temperature and the level of mercury in a thermometer will disappear if we don't hold fixed a good deal of local infrastructure. These connections are contingent on that local infrastructure and don't hold generally. All of the interesting structure is contained in the local sub-models, and emerges relative to constraints and holding fixed certain elements of internal structure.

So the *direction* of influence is introduced by the choice of exogenous and endogenous variables, and the *strength* of the connection depends on

[17] See Price (2007).

the specification of invariants. The upshot is that the idea that changes in one variable *bring about* changes in another is imposed by these choices. It's not an internal relation between the events. The direction in which influence between A and B runs depends on which you treat as exogenous and which you treat as endogenous. This doesn't mean that there aren't objective facts about what varies with what *relative to a choice of sub-model*. It just means that we have to make the choice of exogenous variables and invariants explicit when we raise such questions. There *are* facts about what varies *under hypothetical variation in X when Y is held fixed*. There is no fact about what is *really* or *absolutely* fixed, or what is *really* or *absolutely* variable. The variation in question is in every case hypothetical, and what varies with it is always relative to a choice of invariants or range of applicability.[18] This is all that realism about causal structure requires. In many ways, it's an entirely commonsensical picture that is very natural for the people with hands-on experience investigating causal structure.

The Pearl Inversion: A reorientation in our understanding of modality

I've gone into this in some detail because it is important to see that the idea that the wide-scope models replace, subsume, or reduce the narrow-scope sub-models that they embed is a mistake about the logic of these relations.[19] Wide-scope models don't override, displace, or compete with narrow-scope models. Models that draw the line between fixed and variable structure in different places complement one another, revealing different aspects of the world's modal substructure. The narrow scope models contain causal information that the wide-scope models simply leave out. The rules that govern the behavior of the parts of a system independently are *modally*

[18] Again, it is worth emphasizing that the invariants are usually specified indirectly by fixing a range of applicability and *discovering* what remains fixed across that range. This is why causal structure is a discovery rather than a stipulation. Perhaps the right way to say it is that if we fix the range of applicability, the world will fix the invariants. But we can always change the range of applicability, and when we do, the modalities will change as well.

[19] What makes the discovery of causal structure possible, where it is possible, is the existence of processes that have the effect of randomizing the values of exogenous variables. We can 'carve off a piece of the world', and actually study the effects of random variation in exogenous variables. The results obtained will be only as good as our capacity to randomize and there are complicated practical issues about how to achieve and ensure randomization, but there is nothing suspicious about the status of this extra modal information.

stronger than the rules that govern them in configuration in the sense that they cover situations that don't arise in configuration. They have counterfactual implications that outrun the counterfactual implications of the rules that apply to the configuration. So far, this is just an observation about the logic of the relations among sub-models, that is, an artifact of the embedding relation, a consequence of the fact that when you embed a narrow-scope model in a wider scope one in a manner that constrains the values of variables that are treated as free in the former, you lose modal information.

But it carries an important lesson. There is an unexamined presumption in the philosophical literature that any modal truths must be derivative of global laws. But on looking through the literature, I could find no good argument for this presumption. There's a mereological principle that all of the *categorical* facts in a narrow-scope model are included in a wide-scope model that embeds it. But I see no modal analogue. And reflecting on the lessons from the engine militate against it. If we used to presume that the rules that govern the parts of a complex system must be derivative of rules that govern the whole, we now see that the order of priority runs in the other direction. Rules that govern the whole are derivative of rules that govern the parts. We start with the basic building blocks with a great deal of freedom of movement and build up more complex systems by restricting their relative motion. A full account of the modal substructure of a complex system would have to recover the rules that pertain to the parts as well as the whole.[20] This is Pearl's attitude announced in the preface to causality, as he says:

> [I used to think that] causality simply provides useful ways of abbreviating and organizing intricate patterns of probabilistic relationships. Today, my view is quite different. I now take causal relationships to be the fundamental building blocks both of physical reality and of human understanding of that reality, and I regard probabilistic relationships as but the surface phenomena of the causal machinery that underlies and propels our understanding of the world. (2000: xiii–xiv)

In my mind, this is a very deep insight. The struggle to derive causal information from global dynamical laws prompted the reflection that led to it, but it is a conclusion that has consequences for how we think about modality quite generally. Realism about causal structure is realism about

[20] The parts at different levels of decomposition and relative to different classes of invariants. It would allow the assessment of intervention counterfactuals for any choice of exogenous and endogenous variables, relative to any class of invariants.

the modal substructure that underlies regularities in the actual pattern of events. This modal substructure underwrites, but is much richer than, the modalities embodied in global laws.

Let's recap, separating what we say about the logic of causal claims from any claim about the metaphysics. We say that causal information is information about rules that describe the behavior of the parts of the system, individually and in configuration, where 'rule' is being used here in a neutral, formal way, to mean 'counterfactual supporting regularity'. Rules that pertain to the system as a whole, 'in full assembly', are stronger than those that pertain to the parts individually. The causal structure of the system captures all of the information about the rules that describe the behavior of the parts individually and in configuration. Causal realism is realism about the relativized counterfactuals: what would happen to A under free variation of B, holding fixed a specified class of invariants? The local rules (rules pertaining to open subsystems interacting with an environment)[21] are prior. The global rules are derivative, obtained by adding constraints on configuration of parts. Experimentally, we discern causal structure by studying the parts separately and then in configuration, holding fixed different aspects of internal structure and external scaffolding. Regularities of various kinds can provide clues to causal structure, but the possibility of intervention and controlled experiment is essential.[22]

If there is no conflict between sub-models that draw the line between fixed and variable structure in different places, how do we choose which sub-models to employ in situ? The choice is problem driven, and context dependent. Practical considerations define the scope of a model. We construct sub-models for purposes at hand, choosing what to include, what to hold fixed and what to vary in a manner dictated by the problem context. In a laboratory setting, we're usually dealing with an object system and treating a collection of identified variables whose values we are interested in and have ways of controlling as exogenous. When we are

[21] There is a little terminological ambiguity here. I follow practice in the interventionist literature of calling systems 'open systems', meaning systems that are subject to external action. In thermodynamics, the notion is used in a slightly different manner, but it won't matter too much here.

[22] I haven't said anything about the metaphysics, but we get a prima facie intuitive understanding from the engine example of how modal facts can be grounded in the concrete material configuration of parts of which a system is composed, each with its own intrinsic range of motion, but bound into a fixed configuration across a range of circumstances.

trying to solve a decision problem we treat our actions as free variables and assess the expected outcomes of different choices. But we model ourselves, too, as individuals and as a group, treating our own actions as endogenous variables and seeing how they respond to differences in others. Recognizing these kinds of patterns can help us understand the cultural, environmental, and sub-personal forces that shape our own behavior.

In decision contexts, what we hold fixed is a partly causal question that depends on what we can expect to be fixed in the hypothetical circumstances in which the choice will take place. And this brings out why decisions are hard. If I am wondering whether I should move to Miami for part of the fall, I can't just hold everything that is the case now fixed, I have to know how the weather will be at the time and whether my children will still be in school. And I have to judge how my priorities and feelings will have evolved in the meantime. Judging that I will want to eat in an hour, even though I'm not hungry now, is a relatively easy call, but judging whether I'll want to be with a partner twenty years down the road, or how I'll feel about having children if I take the plunge now, are not. There is no simple recipe for making these judgments. They are causal judgments, but ones that demand self-understanding and practical wisdom beyond mere scientific knowledge. And this is to say nothing about all of the hard moral questions. What do we hold fixed for the purposes of assessing responsibility for events? What should we hold fixed and allow to vary for the purposes of assigning praise and blame? What should we blame for a car accident that injures a pedestrian? Do we blame the driver's slow reflexes? Those caused the crash only if we hold fixed the weather, the worn tread on the tire, the other driver's poor eyesight and the slow pace of the pedestrian? Who is to be blamed for the Gulf oil spill or the Japanese nuclear disaster? In systems that contain multiple interacting human and non-human components, there is no absolute answer to the question 'who, or what, is to blame?' Where the locus of control lies will depend on what we hold fixed and what we allow to vary.

The problem of causal antecedents defused?

Let's revisit the problem of causal antecedents with this understanding of the deep logic of causal judgments in hand. To get the purported conflict with free will going we are invited to see action in the context of wider embedding models—psychological models, neuroscientific models, sociocultural models—in which action appears constrained by exogenous variables. The worry about physical determinism is the most extreme version of this sort of model in which action is not just *constrained* but *determined*, and determined by variables that we can trace to the beginning of time. These strategies for undermining freedom purport to take the reigns out of our hands. The wide-scope view of action they provide is supposed to override the more limited view we adopt for decision and show us that treating our own actions as free variables is deluded. But that is a mistake. There is no incompatibility between the wide-scope and narrow-scope views of action. Our actions appear in multiple models, sometimes as exogenous, sometimes as endogenous. Whenever we model a subsystem of the world—whether it is a car engine, or a human being—we can always widen our view to attain a more encompassing perspective, and wide-scope models will have a set of possibilities that is typically constrained relative to the narrow scope model.[23] But there is no more conflict between these models than there is between the view of a building from close-up and the view from a very great distance.

The choice between models is pragmatic rather than descriptive. Where our choices make a difference to what happens, we use models that treat our choices as free variables so that our choices can be guided by their projected outcomes.[24] In that context, the hypothetical possibilities that matter are the ones that correspond to different choices and hold fixed features of the world that are invariant under choice and expected to obtain in the context of action.

Science, metaphysics, and common sense

I have been suggesting that the concept of causation formalized in work on the logic of causal modeling defuses the apparent conflict between the first personal view of action in decision contexts and the wide-scope view of action as it appears *sub specie aeternitatus*, making the choice between models a pragmatic one and defusing the idea that causes exercise a freedom-undermining compulsion over their effects. Causal thinking is a tool, a

[23] The only exception is the degenerate case in which the newly exogenous variables are random relative to the old across the relevant range of applicability.

[24] For more on the logic of the choice context, see my 'Decision and the Open Future', forthcoming.

cognitive strategy that exploits the network of relatively fixed connections among the components of a system for strategic action, *not* a coercive force built into the fabric of nature that imbues some events with the power to bring about others. Far from undermining freedom, the existence of causal pathways is what creates space for the emergence of will because it opens up the space for effective choice.

There is an ongoing discussion of causal notions in analytic metaphysics that proceeds by eliciting intuitions about hypothetical scenarios in an attempt to systematize everyday intuitions about when it is right to say that A caused (or was causally implicated in the occurrence of) B. Typical discussions start with a paradigm example (e.g. Suzy throws a rock that causes a window to break). An analysis is offered (e.g. c is a cause of e exactly if e wouldn't have occurred if c hadn't occurred). Then problem cases and apparent counterexamples are entertained to try to refine the analysis (e.g. pre-emptive causes, preventive causes, the transitivity of causation, and overdetermination).[25] A charitable construal of the goal of these discussions is to provide an analysis of the everyday notion of cause, or to explicate the folk theory that underlies everyday causal judgments. The methodology and the goal in what I've been doing are both very different. There is no consulting intuitions or armchair reflection on the concept of cause. And what is offered is not an *analysis* of the everyday notion. It doesn't purport to give the meaning of the word 'cause', or the content of anybody's causal concept. It is, rather, as an explication of the scientific notion of cause, which is a refinement and generalization of everyday causal notions. The everyday concept of cause is a clear ancestor, reconstructed in retrospect as a crude, folk version of the developed notions, less precise and less general because conditioned on contingencies about our selves and the context of use.

Science is in the business of giving an account of the universe and our place in it, and to that extent, it is a form of metaphysical inquiry. How should someone looking to science for answers to metaphysical questions proceed? Let the science develop and see where it leads. Physics seems to speak of matters that are thoroughly familiar. But as it increases the scope and depth of its descriptions, it transforms and generalizes those notions. Think of how our concepts of light, matter, and sound have been transformed by science, to say nothing of space, time, and motion. That is what

[25] See Collins, Hall, and Paul (2004) for a good sample of this literature.

has happened with cause. The everyday notion of an intrinsic, compulsive force between natural events has been quite thoroughly transformed. What science tells us is really there—or if you like, 'there anyway'—at the fundamental level as modal substructure is not intrinsically directed, and not an internal relation between pairs of events. The structure captured in causal models identifies relations among variables always relative to some class of invariants, with a direction imposed by choices of endogenous and exogenous variables.

The refinement and generalization that causal notions undergo in the hands of science is characteristic of the scientific development of everyday notions. It replaces the ideas of common sense with more exact and general notions. Pragmatic, contextual sensitivity is eliminated in favor of explicit relativization. Features of the pre-theoretic concept that have an analytic or *a priori* character are often revealed as conditioned on contingencies from a more fundamental perspective.[26] The transformations tend to be conservative of the important core of our practices with the notions, but they can be ruthless with the metaphysical pictures that often accompany those practices.[27]

I don't doubt that there are conceptions of causal notions in the metaphysics literature that make it out to be the kind of asymmetric compulsive force writ into the fabric of nature that would undermine personal freedom, and even that such conceptions may come closer to capturing the core of the folk conception of cause. But the naturalistic metaphysician takes her concept of cause from science, and it is not fair to impose prescientific ideas of causation in an argument that science makes no room for free will. Not only is there room for personal freedom in a fully articulated naturalistic picture of the world, the naturalistic picture gives us an understanding of ourselves and our place in nature that is much richer and more faithful to the experience of agency than the naïve metaphysical images that accompany common sense.

[26] In this case, the direction of influence turns out to be imposed by the choices of endogenous and exogenous variables, and the specification of invariants that is usually left tacit in the everyday application of those notions (in the form of ill-defined *ceteris paribus* clauses) is made explicit and systematized.

[27] In this case, the practices we want to preserve are the practice of treating choice as a free variable in decision, and that of holding ourselves and other responsible for our choices. I have focused on the first. The second takes separate argument.

Of course, a full naturalistic resolution of the conflict between freedom and causal determinism will require parallel developments in our understanding of freedom. Dennett has been the most outspoken explicit naturalistic revisionist about the common sense notion of free will and *Freedom Evolves* (2003) is a masterful exercise in the style of naturalistic metaphysics that I am recommending. Galen Strawson has complained that Dennett's view doesn't vindicate pre-theoretic conceptions of freedom. He writes:

> [The kind of freedom that Dennett argues we have] . . . doesn't give us what we want and are sure we have: ultimate, buck-stopping responsibility for what we do. It allows, after all, that the whole course of our lives may be fixed down to the last detail before we've even been conceived. (Strawson, 2003: 2)

He may be right. People not only have ideas about causation that give it the kind of coercive unidirectional force to undermine their sense of freedom, but ideas about what freedom amounts to—'ultimate, buck-stopping responsibility'—that conflict with the existence of wide-scope models in which our actions appear causally determined by antecedent variables. But Dennett's reaction to Strawson's complaint—'so much the worse for pre-theoretic conceptions'[28]—is, from the point of view of the style of naturalistic metaphysics that I am recommending, exactly appropriate. The theory of general relativity doesn't preserve pre-theoretic intuitions about space and time, but so much the worse for those intuitions. Someone looking to preserve pre-theoretic ideas will find little satisfaction in science. But someone looking to explore what kind of freedom a naturalistic picture of ourselves allows, will find reinforcement of her practices of choosing and valuing, placed on stronger foundations and with more explicit positive criteria for making informed choices.

Connections

There is little to disagree with from my perspective in the rest of *Freedom Evolves* but the chapter on causation could be strengthened in a way that takes advantage of formalism and developments that Dennett didn't have a chance to exploit. In the 'Notes on Sources and Further Reading', Dennett remarks that he discovered Pearl's book only while preparing the final draft, saying:

[28] These are my words, not his, but a fair paraphrase.

> Judea Pearl's *Causality* . . . raises questions about [my way of putting things], while opening up tempting alternative accounts. It will be no small labor to digest these and, if need be, reformulate our conclusions . . . this is work for the future. (2003: 95)

It's work that I've tried to do here. The discussion can also, however, serve as a model of a style of naturalistic metaphysics that picks up some of the themes in *Every Thing Must Go*. One of the complaints that critics have made about the book is that they don't provide a clear positive alternative to analytic metaphysics; a clear idea of what naturalistic metaphysics might be like. Cian Dorr, for example, writes in *Notre Dame Philosophical Reviews*:

> The authors' self-proclaimedly 'scientistic' view of the supremacy of the sciences among forms of human inquiry makes it prima facie puzzling how there could be room for a field called 'metaphysics' that is not itself a science and is nevertheless not a waste of time. One traditional picture imagines metaphysics as standing to physics as physics stands to biology—investigating a proprietary realm of facts that cannot even be expressed in the vocabulary of physics, and that provide 'foundations' for physics in the same sense in which physics provides foundations for biology. The authors are skeptical about whether there is any such realm of facts, and thus whether there is any legitimate enterprise of investigating them. What else might 'metaphysics' be, then? (Dorr, 2010)

I certainly agree with Ladyman and Ross that there is no science-independent standpoint from which to do metaphysics. I don't know what naturalistic metaphysics is if it is not simply to bring everything we learn from science—about the world and about our place in it—to bear on philosophical problems, and to grant the scientific view authority over any pre-theoretic intuitions we might have about how things are. I have tried to give concrete content to one way of doing that here, but I don't think there is likely to be any simple template for naturalistic metaphysics or any simple way of demarcating it from science.

Even this weak view of naturalistic metaphysics has opponents, if this closing remark from Dorr's review is anything to go by:

> *Every Thing Must Go* seems likely, if it has any effect at all on analytic metaphysicians, merely to confirm a few more of them in their impression that no one has yet shown how developments in the sciences might be relevant to their concerns. (2010, last line of review)

I'm not sure whom he's speaking of or whether he counts himself in that camp. Most of the analytic metaphysicians I know are engaged with and interested in developments in science. But if his own remark has any effect it will most likely be to confirm philosophers of science in an increasingly pronounced distaste for unscientific metaphysics.

References

Bohm, D. (1989) *Quantum Theory*. New York: Dover.

Born, N. (1971). *The Born-Einstein Letters*. New York: Walter.

Cartwright, N. (1979). Causal laws and effective strategies. *Noûs* 13: 419–38.

Hall, N., and Paul, L.(eds.) (2004). *Causation and Counterfactuals*. Cambridge MA: MIT Press.

Dennett, D. (2003). *Freedom Evolves*. New York: Viking Press.

Dorr, C. (2010). Review of *Every Thing Must Go: Metaphysics Naturalized*. *Notre Dame Philosophical Reviews* [online]. Available: <http://ndpr.nd.edu/news/24377/?id=19947> accessed 17 May 2012.

Ferraiolo, W. (2004). Against compatibilism: Compulsion, free agency and moral responsibility. *Sorites*, 15: 67–72.

Ismael, J. (forthcoming). Decision and the open future. In A. Bardon, ed., *The Future of the Philosophy of Time*. Oxford: Oxford University Press.

Ladyman, J., and Ross, D. (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.

Pearl, J. (2000). *Causality*. Oxford: Oxford University Press.

Price, H. (2007). Causal perspectivalism. In H. Price and R. Corry. eds., *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*. Oxford: Oxford University Press.

Russell, B. (1918). On the notion of cause. In *Mysticism and Logic*. London: Allen & Unwin.

Strawson, G. (2003). Evolution explains it all for you. *New York Times* [online]. Available: <http://www.nytimes.com/2003/03/02/books/evolution-explains-it-all-for-you.html?src=pm> accessed 17 May 2012.

Taylor, R. (1963). *Metaphysics*. Englewood Cliffs, NJ: Prentice-Hall,

Woodward, J. (2003). *Making Things Happen*. Oxford: Oxford University Press.