6

CAUSAL CONTENT AND GLOBAL LAWS

Grounding Modality in Experimental Practice

JENANN ISMAEL

When philosophers approach science, there is a strong tendency to focus on the products of science rather than the practice. What follows is a case study in how that focus has distorted the philosophical discussion of science and laws, and how an emphasis on the experimental side of modeling resolves them. The focus on the products of science leads to a fixation on global models and a tendency to treat them representationally. An emphasis on the experimental side of modeling, by contrast, directs our gaze toward models of open subsystems and gives us the tools for a more pragmatic approach to modal content.

There was a time when science was thought of almost exclusively in causal terms. The mandate of science was thought to be the investigation of the causal structure of the world. Things changed with the mathematicization of science and the triumph of Newtonian theory. Newton's theory provided dynamical laws expressed in the form of differential equations that could be used to compute the state of the world at one time as a function of its state at another. Theoretical developments since Newton have seen important departures. Quantum mechanics introduced indeterminism. We are still far from a fundamental theory, but in the philosophical literature, Newtonian theory still serves in many circles as a paradigm for what a fundamental theory should look like: it should be global in scope, and the fundamental laws should take the form of what we call equations of motion, which is to say they tell us how the state of the world evolves from one moment to the next.

The eclipsing of causal notions in physics happened almost unnoticed until Russell's justly famous paper of 1913 in which he detailed the differences between the two notions (Russell 1913/1953). Russell himself thought that the differences were so great, and that the causal notions carried so many ill-fitting associations, that they should be eliminated from exact science. That early position was rebutted in a paper by Nancy Cartwright (1979) in which she argued that causal knowledge is indispensable in practical reasoning. And so began the long struggle to understand how causal ideas enter into scientific description.[1]

REDUCTIVE PROJECTS

Reductive projects dominate the rather large body of post-Russellian discussion of causation, especially among philosophers of physics. The thought is that what happened with cause is what happened with so many other notions once thought to be basic to our understanding of nature. Once the physics has progressed so that cause no longer makes an appearance at the fundamental level, there is room for an illuminating reduction, so we should be looking to reduce causal structure to physical laws. And the presumption was that these laws should take the form of global laws of temporal evolution, modeled on the Newtonian law of gravitation. These were the kinds of laws that Russell regarded as the most basic modal generalizations in physics. He wrote,

> In the motions of mutually gravitating bodies, there is nothing that can be called a cause, and nothing that can be called an effect; there is merely a formula. Certain differential equations can be found, which hold at every instant for every particle of the system, and which, given the configuration and velocities at one instant, or the configurations at two instants, render the configuration at any other earlier or later instant theoretically calculable. That is to say, the configuration at any instant is a function of that instant and the configurations at two given instants. (Russell 1913/1953, 14)

Let us call the assumption that the most basic laws take the form of global laws of temporal evolution—that is, laws that give the state of the universe at one time as a function of its state at another—*globalism*. Most philosophers have followed Russell in presuming globalism. I say "presuming" because globalism is almost without exception assumed without argument, and it plays a role shaping the conception of physical necessity both inside and outside philosophy of science.[2]

I want to use the post-Russellian discussion of causation to argue that globalism is subtly but importantly mistaken. Here is my discussion plan: first I will say a little about why causal structure is not reducible to global laws, and then I will make a case for causal realism. After that, I will suggest that global laws are derivative from local laws that take the form of (what I will call) rules for mechanisms. Then, in the next section, I will argue that physics was never globalist, and finally I will talk about how rejecting globalism reopens the possibility of grounding causal structure in fundamental laws. Perhaps the most important upshot of the discussion is that it paves the way for an empiricist account of physical necessity, one that can be connected quite directly to experimental practice. I will conclude with some remarks about this.

THE CONTENT OF CAUSAL CLAIMS

Understanding the content of causal judgments has been a long, hard, and heavily contested road. For some time the philosophical discussion was dominated by attempts to provide analyses that systematize everyday intuitions about when *A* causes *B*. In recent years, something of a revolution occurred led by developments in cognitive and computer science, psychology, and statistics.[3] Instead of trying to systematize everyday intuitions about causes, attention turned to providing a formal framework for representing causal relations in science. We now have such a framework in the interventionist account of the content of causal claims that also sheds a good deal of light on everyday causal claims. The interventionist account came out of independent work by Glymour's group at Carnegie Mellon and Judea Pearl at the University of California–Los Angeles. Pearl's work culminated in his *Causality* (2000), though many people in the philosophical literature know interventionism from Woodward's *Making Things Happen* (2003a). I rely here on Pearl.

Pearl is a computer scientist and statistician, and he approached discussion of causal structure initially with the Bayesian presumptions that dominate his field. He writes:

> In order to be combined with data, our knowledge must first be cast in some formal language, and what I have come to realize in the past ten years is that the language of probability is not suitable for the task; the bulk of human knowledge is organized around causal, not probabilistic relationships, and the grammar of probability calculus is insufficient for capturing those relation-

> ships. Specifically, the building blocks of our scientific and everyday knowledge are elementary facts such as "mud does not cause rain" and "symptoms do not cause disease" and those facts, strangely enough, cannot be expressed in the vocabulary of probability calculus. (2001, 19)

He set out to do for causal information what the probability calculus does for probabilities, supplementing the probability calculus with a formalism that was adequate to the expression of causal information. Pearl's goal was, in his words, "The enrichment of personal probabilities with causal vocabulary and causal calculus, so as to bring mathematical analysis closer to where knowledge resides" (2001, 19).

In Pearl's account, causes are to practical reasoning what probabilities are to epistemic reasoning. Whereas probabilities provide information about the correlations among a collection V of variables, causal information adds counterfactual information about *how changes in the value of one variable induce changes in the value of others.* Singular causal claims depend on generic causal information. Generic causal information is information about how one variable in a network induces changes in the values of other variables. Before we have a well-defined question about whether X_i is a cause of X_j, we have to specify a network. Once the network is specified, the question of whether X_i is a cause of X_j is the question of whether interventions on X_i induce changes in the value of X_j. An intervention on X_i is a change in its values that is "surgical" in the sense that it severs the connection between X_i and its parents in the network.[4] So if X_i is one of the variables in the network formed by V, knowing the causal effects of X_i is knowing what would happen if X_i were separated out of this web, severing connections with its own past causes and allowing it to vary. An "intervention" is just a formal name for the virtual act of separating a variable from its past causes.

Direct causation, represented by an arrow, is the most basic causal relation. A variable X_i is a direct cause of another variable X_j, relative to a variable set V, just in case there is an intervention on X_i that will change the value of X_j (or the probability distribution over the values of X_j) when all variables in V except X_i and X_j are held fixed (Pearl, 2001, 55). Causal relations are relative to networks. As variables are added to the network, new arrows appear, and others disappear. Causal relations of many kinds including total, direct, and indirect causes, necessary and sufficient causes, and actual and generic causes have been defined and successfully analyzed within the formal framework of structural causal models (SCM).

Scientists and philosophers of science have seized on Pearl's formalism for capturing the causal content of science. Psychologists and social scientists and econometricians have put it to work. The SCM framework is not an analysis of the folk concept of cause.[5] What it does rather is systematize the patterns of counterfactual judgments conveyed by everyday causal statements and needed to play the role identified by Cartwright in practical reasoning. What made Russell's eliminativism an insupportable position was not that science has to preserve folk intuitions about the world, but that causal reasoning and practical reasoning go hand in hand. Causal judgments supply the counterfactuals needed to identify strategic routes to action. When one is choosing between alternative courses of action, one needs to assess counterfactuals of the form "what *would* happen if I *A*'ed (rather than *B*'ed or *C*'ed)?" The builder choosing between wood and steel needs to know what would happen if he chose wood *and* what would happen if he chose steel, even though he will only do one or the other. The traveler choosing between path A and path B needs to know what would happen if he chose A *and* what would happen if he chose B, even though he cannot travel both. SCM shows us how to express claims of this form in exact logical terms and shows us why those claims are not captured by probabilities. As Pearl says,

> Probability theory deals with beliefs about an uncertain, yet static world, while causality deals with changes that occur in the world itself (or in one's theory of such changes). More specifically, causality deals with how probability functions change in response to influences (e.g., new conditions or interventions) that originate from outside the probability space, while probability theory, even when given a fully specified joint density function on all (temporally-indexed) variables in the space, cannot tell us how that function would change under such external influences. (2001, 36)

The interventionist analysis holds an important lesson for the proponents of the reductionist project. There is a very simple logical point that can be extracted from it that explains why the counterfactuals that we need to play the role of causal beliefs in deliberative reasoning will not in general be extractable from global dynamical laws. Let S be a system governed by a dynamical law L that gives the state of S at any time as a function of its state at earlier times. Let us suppose that it is a consequence of L that at any given time $A = kB$ (where k is a constant), and that there is some variable C downstream of A and B such that $C = f(A, B)$ (or more explicitly $C_t = f(A_{t-n}, B_{t-n})$, but I will suppress the temporal parameters).[6] There are three separate causal hypotheses not discriminated by the information that is given:

1. *A* causes *C.*
2. *B* causes *C.*
3. Neither *A* nor *B* cause *C. A, B,* and *C* are joint effects of a cause *D* in their common past.

Hypotheses 1, 2, and 3 are distinguished by the pattern of counterfactual dependence they postulate among *A, B,* and *C.* Hypothesis 1 entails that if *B* is held fixed and *A* is allowed to vary, *C* would vary as well. Hypotheses 2 and 3 entail that it would not. Hypothesis 2 entails that if *A* is held fixed and *B* is allowed to vary, *C* would vary as well. Hypotheses 1 and 3 entail that it would not. Hypothesis 3 entails that if *A* and *B* were held fixed and *D* were allowed to vary, *C* would vary as well. Hypotheses 1 and 2 entail that it would not.

The reason that the dynamical laws do not discriminate among these hypotheses is that the antecedents are counter*legals:* there are no models of the global laws in which the values of *A* and *B* vary independently.[7] In all of those models at all times, $A = kB$. Global laws underdetermine patterns of counterfactual dependence at the local level whenever there are interventions whose antecedents are not nomologically possible.[8] The information contained in a causal model is, in general, *strictly logically stronger* than the information contained in the global laws. Any law-like constraint on coevolution of local parameters is going to be preserved by evolution, and only the result of hypothetical interventions whose antecedents are counterlegals is going to separate the causal hypotheses.

THE CASE FOR CAUSAL REALISM

Let us pause to take stock. Russell observed that causal relations do not appear in a fundamental theory. He suggested that the notion of cause is a folk notion that has been superseded by global laws of temporal evolution and has no place in exact science. Cartwright observed that causal information plays an indispensable role in practical reasoning—that is, that the functional essence of causal beliefs is to supply the information about the results of hypothetical interventions needed for practical reasoning. This was codified in the interventionist analysis, which provided a precise formal framework for representing and investigating causal relationships,[9] and made it easy to see why causal information outruns the information generally contained in global dynamical laws.

To react to this situation, it will be useful to have a better intuitive feel for the relationship between causal facts and global laws. The real beauty of the SCM is that it gives insight into what grounds modal claims in science. Causal models are generalizations of the structural equations used in engineering, biology, economics, and social science. In a causal model, a complex system is represented as a modular collection of stable and autonomous components called "mechanisms." The behavior of each of these is represented as a function, and changes due to interventions are treated as local modifications of these functions. The dynamical law for the whole is recovered by assembling these in a configuration that imposes constraints on their relative variation. If we know how a complex system decomposes into mechanism, we know how interventions on the input to one mechanism propagate through the system. But since there are many ways of putting together mechanisms to get the same evolution at the global level, we cannot in general recover the causal information from the global dynamics.

Consider a complex mechanical system like a washing machine. We can model such a machine as a unit and write down an equation that allows us to calculate its state at one time from its state at any other: $[S_{\text{final}} = f(S_{\text{initial}})]$. If we were simply interested in description or prediction, and we knew the initial state of the engine, this would tell us everything there was to know. But if we want the kind of working knowledge that would let us troubleshoot, or intercede to modify the machine's behavior, *f* would not be enough. That kind of information is usually conveyed by a diagram that decomposes the machine into separable components, indicates how the components would behave if the parts were separated out and their input allowed to vary without constraint, and helps the viewer form an understanding of how the fixed connections among the parts within the context of the machine produce the overall pattern of behavior.[10] The global dynamics for the machine as a whole contain information about how its state varies over time, but it is the decomposition into mechanisms that tells us what would happen under interventions that do not arise under the normal evolution of the assembled machine.

In logical terms, the reason causal information goes missing when we just give the dynamics for the overall state is that the rules governing components are modally richer than the rule governing the whole. Embedded in a larger machine, the input to each component is constrained by the fixed connections among the parts of the engine so that each now moves within a restricted range and the patterns of counterfactual dependence that the interventionist sees as crucial to causal claims are lost. Variables that were allowed to vary freely in the original model are constrained by the values of variables in the embedding model, so information about what *would* happen *if* they were allowed to vary without constraint is lost. The dynamical law for the whole is not *incompatible* with the laws that govern the components. It is just that if we only look at the way that a system evolves as a whole, we lose information about the modal substructure. There is a real and absolute loss of modal information that occurs when one moves from a narrow-scope model of a subsystem to a wider scope model in which the input to the submodel is constrained by the values of variables included in the wider model. Different ways of piecing together subsystems, with different implications for the results of local interventions, preserve the global dynamics.

REACTION: THE PEARL INVERSION

One might react to this situation by denying that there are modal facts over and above those that can be derived from global laws. In this view, one says that there is no fact of the matter about what would happen if some parameter were magically separated out and allowed to vary freely. There are two things to say about this. The first recapitulates Cartwright's response to Russellian eliminativism. Causal information is indispensable in practical reasoning. If we were simply interested in prediction, laws for the cosmos combined with information about initial conditions would tell us everything we need to know. But we face choices about how to act in the world, and we are interested in knowing how various ways of acting *would* play out. This is a way of saying that our actions are nodes in causal networks. Whenever I ask "what would happen if I *A*'ed rather the *B*'ed," I am asking for specifically causal information about the effects of *A*'ing. We care about causal information because we are not mere observers of nature but agents, and our actions have for us the status of interventions.[11] For purposes of predicting whether an engine will break down, it does not matter whether dirty oil causes or is merely a sign of impending engine breakdown.[12] It does not matter, that is to say, whether dirty oil is the symptom or the disease. The causal information matters for the mechanic who needs to know whether cleaning the oil will solve the problem.[13] We do not encounter the world as observers. We encounter it as agents. To act on the world, we need to know how it would respond to potential interventions.

The second thing to say questions the motive for rejecting causal information in favor of global laws. The components of the world are open systems. We encounter them severally, in multiple settings. Ideally we can isolate them in the laboratory and study their behavior individually and in interaction with other systems. Do we really think that the counterfactual implications of global laws purporting to describe completely specified alternatives to actuality are on a more secure epistemic or conceptual footing than the counterfactual implications of models of open subsystems of the actual world that we can isolate in the laboratory and study under conditions that approximate intervention?[14] Philosophers tend to be uninterested in partial views of bits of the world. They make a lunge for the most encompassing view. Most day-to-day science, however, is not concerned with the world as a unit, but is focused on local subsystems. The experimental scientist does his best to carve off a manageable bit of the universe. In the best case, his study is more or less tightly focused on a smaller unit, which can be isolated in the laboratory and whose responses to controlled interventions can be observed. That is not possible with larger systems, but we piece together an understanding of larger systems from an understanding of the rules governing components in constrained configuration.[15]

Models of open subsystems do have modal content: they identify counterfactual supporting regularities, so they involve induction from the observed results of actual interventions to merely potential ones. But the modal content is empirically grounded in testable regularities. When developing a model of an open subsystem, the scientist isolates the system as well as she can, identifies the variables whose causal effects she is interested in, finds some way of manipulating them while holding fixed the features of the internal configuration and environment she is imposing as constraints, and observes the effects. While there are practical difficulties in experimentally realizing situations to test for particular modal claims, there is nothing in principle untestable about modal claims pertaining to open subsystems of the universe.[16]

Things are different at the global level. Modal claims at that level purport to describe completely specified alternatives to actuality. They involve inductions from one case (the actual world) to a space of merely possible worlds. Possible worlds are entirely extrinsic to the actual world and are not even potentially observed. It is hard to say not only how we know about non-actual possible worlds but also why we care about them. The question of how an open system would behave if acted on in various ways, by contrast, has a transparent connection to actual things and an obvious practical interest. The question of how an open system would behave if acted on in various ways is what gives the idea of modality its practical significance. This idea gets generalized in causal models and applied to hypothetical interventions that go beyond what we can actually effect, but without losing its significance (in the same way that the idea of a view of the world from positions in space that we have no way of getting to makes good, if conjectural, sense). We get a problem if we reverse the order of explanation, reduce the local claims to global ones, then struggle to find an interpretation for these purportedly global modal facts. Unlike modal claims pertaining to open subsystems of the universe, which can be understood in terms of how those systems would respond to hypothetical interventions, the modal content of global models is *strictly* and *irremediably* counterfactual. Its semantic content is strictly extrinsic to the actual world.

In practice, we arrive at global laws like Newton's by extrapolation from the laws that govern its components. The modal implications of the global form of that law have no empirical or practical significance of their own. We are able to form beliefs about what would happen under hypothetical conditions because the world is composed of mechanisms that can be investigated independently and then recombined into larger systems whose behavior is a function of the rules governing components. The modal content trickles up from the experimentally based understanding of relatively simple components to larger configurations, rather than the other way around. It begins with modal generalizations that apply to the sorts of controlled subsystems that we can study in the laboratory. We form an understanding of larger units by piecing together what we know of the components: how they behave individually and in interaction when they are allowed to move freely and in constrained configurations. That is how knowledge obtained in the laboratory can lead to empirically well-founded beliefs about configurations that have not themselves been studied, and to empirically well-founded beliefs about configurations that have been studied to interventions that have not been observed. When we construct a new bridge or building, one that is not a copy of anything that has gone before, or when we synthesize a new pharmaceutical agent, we are not making wild inductive leaps of the kind we associate with theoretical breakthroughs. We are combining well-understood components in new ways. Every piece of new technology designed on paper that behaves as expected is possible because our modal knowledge is compositional in this way. When we understand how the components behave and the compositional principles, in a great many cases, we understand configurations.

In practice, the empirical content of laws attaches to predictions for open subsystems. These are derived not from global laws but from the rules pertaining to the mechanisms of which they are composed.

PHYSICS AND NATURE'S ULTIMATE MECHANISMS

One might argue that there is something quite misleading about these examples; we get this result only because in these examples we are tacitly restricting attention to global possibilities that leave the machine configuration intact. And the intuition that there is missing modal information is trading on the fact that we can imagine taking the machine apart and reconfiguring its parts. But if there are no law-like restrictions on the composition of mechanisms, then there is a global possibility for every local intervention and the difference disappears. So if we insist that in a well-behaved theory there is a global possibility that corresponds to every describable reconfiguration of components, then although there will not be an actual intervention that alters any of what we regard as the frozen accidents of our world (i.e., contingent features of initial conditions preserved by temporal evolution), there will be a model of the laws in which the antecedent of any intervention counterfactual holds.[17] Another way to put this is that if we are given the phase space for the universe as a whole and focus attention on the subspace that corresponds to the state of any subsystem of the world, treating everything else as exogenous, in a well-behaved theory (i.e., one in which there are no ultimate restrictions on configuration of components) it is arguable that there is always going to be a possible global state for every point on the boundary of the subspace. And if this is correct, then we can substitute global laws with a combinatorial property and capture the logical content of rules for mechanisms. So long as there is a global possibility for every local intervention, the laws pertaining to mechanisms will be recoverable from global laws.

In my view, this observation simply *reinforces* the point that our modal knowledge is rooted in our understanding of mechanisms. In a well-behaved theory, there is a global possibility for every local intervention because it is our ideas about what nature's basic mechanisms are that drive our ideas about what global configurations there are, rather than the other way around (*vide* Ismael 2013). What *makes* such a theory well behaved is that the global possibilities allow recombination of mechanisms.[18] But the observation does bring out a subtle ambiguity in the notion of law. Globalism is the thesis that the global laws of temporal evolution for our world are the most basic nomic generalizations. Those are laws that tell us how to calculate the state of our universe at one time from its state at other times. But there is another notion of law that is much more general, according to which global laws are laws that tell us not just how the state of our universe varies over time, but what kinds of universes are possible.[19] This is the kind of law that a fundamental theory in physics gives us. When such a theory is presented in physics, laws are given for simple components, and laws for complex systems are built up from those. Any way of piecing components together counts as a possible global configuration. We obtain the law of temporal evolution for our universe by specifying its initial configuration.[20]

Here is where I think we come back to the question of where Russell went wrong. When Russell looked for the most basic physical laws, he took the form of Newton's laws that apply to our world as a whole. What he *should* have done is taken the Newtonian laws governing the basic components of nature—nature's ultimate mechanisms—as basic. The fundamental law of temporal evolution for our world (the one for which Russell used the Newtonian law described in the quoted passage earlier) is the special form that these laws take for a system made up of the particular set of components of which our world is made arranged in a particular way. The composition of our world and its initial configuration encode contingent information needed to obtain detailed predictions. But the modal content of a theory is contained in the laws that govern components and the rules of composition. We cannot in general recover those from the global dynamics. So we have a trade-off in categorical and modal content. The rules for the components are weaker in categorical content than the law of temporal evolution for our world, but richer in modal content. The more categorical content we include, the less modal information we convey. To put it another way, the more information we have about what is actually the case, the less information about what would happen under non-actual conditions.

I remarked earlier that Pearl is not a metaphysician. He approached discussion of causation from the point of view of the statistician. In the preface to *Causality* (2000) he describes how his own thinking shifted away from the Bayesian presumptions that dominated his own field to the view that causal structure was fundamental:

> [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. (xiii–xiv)

I am suggesting a parallel shift away from the globalist presumptions that dominate the philosophy of science. I used to think that talk of mechanisms was a useful way of conveying partial information about global laws. Today, my view is quite different. I now take mechanisms to be the fundamental building blocks both of physical reality and of scientific understanding of that reality, and I regard global laws as but the emergent product of the mechanisms that underlie and propel our understanding of the world.

PHILOSOPHICAL IMPACT

This shift in thinking has several kinds of philosophical impact. In the first place, it reopens the possibility of grounding causal claims in fundamental law, suggesting a rather different research program for a physical fundamentalist. Instead of trying to derive causal facts from global laws, he sees the causal relations captured in directed acyclic graphs (DAGs) as emergent regularities rooted in composition of mechanisms.[21]

In the second place, issues about the metaphysics and epistemology of laws look rather different when rules for mechanisms are substituted for global laws. There is no question that science is steeped in modality. It studies not just what does happen but what could, and must, and would happen under hypothetical conditions. But the modal commitments of science create a dilemma for the empiricist. On one hand, belief in science seems the hallmark of empiricist commitment. On the other hand, believing that the world is governed by global laws, together with the inflated metaphysical commitments which that seems to carry, runs counter to the empiricist instinct.[22]

Directing our hermeneutic attention *away* from global laws toward the kinds of testable regularities pertaining to smaller than world-sized components of nature that we can isolate and study in the laboratory is a positive development for the empiricist. These have a well-behaved epistemology and make the most direct contact with our practical interests.[23] The empiricist can be discriminating about modal claims in science. She should not try to eliminate the modal content of a theory. The modal implications of theory play an indispensable role in guiding our interactions with open subsystems of the world and are grounded in empirically testable regularities. But she can be less tolerant of laws that pertain specifically to worlds as wholes.[24] It is only at the global level that modality becomes weird. To the extent to which they are not mere extrapolations of local modalities, global laws are no longer grounded in testable regularities. They become about other *worlds* rather than a hypothetical variation in our world, and they lose touch with the practical and empirical significance modal claims have for embedded agents. From an empiricist point of view, there is something altogether *upside down* about thinking that modal facts pertaining to the world as a whole are more epistemically or metaphysically secure than modal facts pertaining to open subsystems of the world. To take global laws as primitive and reject modal facts that cannot be reduced to them is to reject something that is immanent, empirically accessible, and metaphysically unmysterious in favor of something that is otherworldly, in principle inaccessible, and metaphysically exotic.[25]

I am strongly inclined to be a realist about modal claims grounded in testable rules for components. The fact that the universe is built up out of mechanisms we can separate from their environments in the laboratory and study in (approximate) isolation is what makes it possible to form modal beliefs. We should not try to reduce the modal implications of our theories, but we should try to ground them in rules for mechanisms, as inductions from testable regularities. I am strongly *dis*inclined, however, to be a realist about modal claims grounded in global laws. Where the globalist says, "Accept modal claims that can be derived from global laws, reject the overflow," I say, "Accept counterfactuals that can be derived from rules for mechanisms, reject the overflow." Or, more cautiously, leaving it open that there might be inductive practices that allow us to make modal inferences at the global level, I say that the burden of proof lies with the person invoking modal claims that cannot be grounded in rules for mechanisms to clarify their empirical basis.[26] What I am deeply suspicious of, however, is a tendency in foundational discussions to invoke global modal claims in an explanatory role. So, for example, in cosmology global laws are invoked to explain why certain global configurations do not arise. But to say that certain kinds of configurations cannot arise because the global laws rule them out strikes me as empty unless the global laws can be derived from some deeper principle.

Finally, this shift in how we think of laws, though in some ways a subtle shift when just thinking about how to express the modal content of science, can have a large impact at the hermeneutic level. Imaginative pictures guide first-order philosophical views, and the imaginative picture that comes with a globalist conception of laws is particularly toxic. The picture of natural necessity as deriving from global laws is very different from one that sees natural necessity as grounded in rules for mechanisms. Instead of ironclad global laws that seem to force history to unfold in one very particular way from its starting point, we have rules that describe the way that nature's simple components behave, something like the rules for chess pieces or the degrees of freedom and ranges of motion that define the behavior of the agitator and drum in the washing machine mentioned previously. These individual rules give rise to complex regularities when the components to which they pertain are placed in different configurations, which can be exploited by well-positioned agents who have control over parts of the machinery (or indeed who *are* parts of the machinery) to bring about more distal ends. The fact that our world is composed of simpler mechanisms that can be isolated and studied in the laboratory is what makes inductive practices and science possible. And that in its turn allows us to identify strategic routes to bringing about ends. It is what allows us to predict and control nature, and to gear our own actions toward desired ends. Laws and causes and all the inductive products of science are part of that. They are not (as the globalist picture encourages us to think) chains that bind us to act as we do. They are handmaids to choice.[27]

This chapter has been advocating a return to a conception of modality grounded in scientific practice and an unwillingness to divorce science from experimental practice. I began with Russell's observation that causes have disappeared from the fundamental level of physical description, gave reasons for thinking that they were nevertheless indispensable for embedded agents, introduced the interventionist account as a formalization of the causal content of science, and showed why the modal content of causal claims generally outruns that of global laws. I then suggested that where Russell went wrong was in taking the global laws of temporal evolution modeled on the law of gravitation for our universe as the most basic nomic generalizations in science; I suggested instead that we take the rules that govern the behavior of nature's basic components as basic.[28]

In making this hermeneutic shift we do several things:

1. We reopen the possibility of recovering at least a large class of causal relationships (those captured in DAGs and formalized by SCM) as emergent regularities grounded in fundamental laws.
2. We clarify the epistemological basis of modal claims in science and pave the way for a moderate empiricist account of alethic modality. Modal judgments in science that can be rooted in rules for mechanisms rather than global laws are just inductions from testable regularities.
3. We free ourselves from an imaginative picture of laws that has played an insidious role outside philosophy of science.[29]

The fixation on global laws is part of a more general tendency among philosophers to focus on the products of science rather than the practice. The best antidote to that tendency is a focus on the experimental side of modeling. If we fixate on global laws, causal structure disappears and becomes difficult to recover. Modal generalizations seem metaphysically mysterious and detached from anything that can be observed. Because experimental practice is by its nature concerned with open subsystems of the world, this directs our gaze away from the global models and toward models of open subsystems. It allows us to connect modal generalizations to testable regularities, established in the laboratory by observing the results of interventions in a controlled setting. And by linking scientific modeling to intervention and manipulation, it gives us the tools for a more pragmatic approach to modal content.

NOTES

1. There are many excellent discussions of Russell's paper and Cartwright's response. See Field (2003) and the papers in Price and Corry (2007). A good deal of the post-Russellian discussion has focused on locating the source of the temporal asymmetry of causation. I will be focusing on a different issue, namely whether the modal content of causal claims can be grounded in fundamental law.

2. See, for example, the notion of laws at work in Helen Steward's discussion of freedom (2012). Her discussion is an unusually explicit but not atypical expression of the notion of physical law that many philosophers take from physics.

3. See Sloman (2005) for a user-friendly summary of these developments.

4. Formally, this amounts to replacing the equation governing X_i with a new equation $X_i = x_i$, substituting for this new value of X_i in all the equations in which X_i occurs, but leaving the other equations themselves unaltered. An intervention is defined as follows: "The simplest type of external intervention is one in which a single variable, say X_i, is forced to take on some fixed value x_i. Such an intervention, which we call 'atomic,' amounts to lifting X_i from the influence of the old functional mechanism $x_i = f_i(pa_i, u_i)$ and placing it under the influence of a new mechanism that sets the value x_i while keeping all other mechanisms unperturbed. Formally, this atomic intervention, which we denote by $do(X_i = x_i)$ or $do(x_i)$ for short, amounts to removing the equation $x_i = f_i(pa_i, u_i)$ from the model and substituting $X_i = x_i$ in the remaining equations" (Pearl, 2000, 70). So for Pearl, once you know what the causal mechanisms are, you can say which interactions constitute interventions. Woodward thinks that this limits the utility of interventions to discover causal mechanisms (among other things) and wants to characterize the notion of an intervention independently so that it can be used as a probe for causal structure. To some extent this in-house dispute reflects a difference in focus. From a metaphysical perspective, it is natural to take the underlying causal structure as basic. It is what explains the surface regularities and patterns of counterfactual dependence. But Woodward is interested in using interventions as a route in, so to speak. He wants to be able to identify interventions (perhaps provisionally) before we have a detailed understanding of the causal structure and use them to probe.

5. Although, see the Appendix to *Causality,* where Pearl has had a lot to say about why causal information is needed for deliberating, and connects it to the working knowledge that we associate with knowledge of how things work. The formalism is linked in this way with everyday notion of cause.

6. This sort of case arises routinely in medical situations in which a doctor needs to distinguish symptoms from cause.

7. One can see this difficulty in the acrobatics that possible worlds' semanticists face assessing counterfactuals that involve local departures from actuality.

8. We can make the same points by talking about phase spaces. When we develop a model of a constrained subsystem of the world, we restrict attention to a subspace of the global phase space. Knowing all the allowed trajectories through the global phase space will not give us the counterfactual

information we need to make causal judgments if there are phase points corresponding to free variation of local variables with no global trajectories through them.

9. Interventionism is an account of the content of causal claims, which we can take to mean an account of the inferential implications of causal beliefs, their role in epistemic and practical reasoning, and their relations to perception and action. For assessment of the SCM formalism and its impact on scientific investigation of cause, see Sloman (2005). For the philosophical development of interventionism and its relationship to alternatives, see Woodward (2003a, 2003b, 321–340).

10. One can look at a different level of resolution, add or subtract variables from the collection, or change background assumptions. Each of these constitutes a change in network and can alter causal relations among nodes. The causal relations at one level of resolution, relative to one collection of variables, and against a given set of background assumptions are different from those at another.

11. See Joyce (2007) and Ismael (2011) for discussion of decisions and their status.

12. The claim is not that it is generally irrelevant but that the specifically causal information does not add anything for predictive purposes to the probabilistic information, because the *specifically causal* information adds only information about unrealized possibilities.

13. Causal knowledge also matters for assigning responsibility for past events and learning from mistakes. It matters for understanding the significance of our choices and gauging their effects. It matters for deciding how to feel about the past and our role in it. Our emotional lives are built around "would have been and could have been."

14. For recent views on the elimination or reduction of causation to physical laws, see Norton (2007) and Maudlin (2007).

15. Sometimes nature creates a natural laboratory in the interaction between an open subsystem and its environment, but that is the exception rather than the rule.

16. There are general, effective ways of isolating a system causally from its environment, shielding it from the effects of exogenous variables. It is possible to operationalize interventions at the local level by inserting a random or pseudo-random process that fixes the value of exogenous variables. In universes in which there were no such processes, it would be impossible to discover causal structure, and science as we know it would be impossible.

17. There will not be interventions that change the energy or charge of a system, for example, just as there will not be interventions that alter accidental correlations among variables preserved by evolution.

18. The combinatorial principle is not a logical truth, and its status is contested. Hypotheses that violate it have irreducibly global constraints on configurations, and there are good reasons for insisting that global constraints should be emergent from rules governing the parts of which the world is composed. This is an issue that needs deeper examination.

19. Compare laws that tell us how a system of a given type evolves over time from any point in its phase space with laws that tell us how to construct the phase spaces for physically possible systems.

20. Where by "initial" we mean only "initial relative to a chosen interval." We do not mean some absolute initial moment in history. Finding the law of temporal evolution for our universe would be akin to specifying the Hamiltonian for our universe.

21. This is an open research program. Carrying it out would involve two components: (1) formalizing the full range of causal claims and supplementing the interventionist framework if it is not adequate to their expression, and (2) investigating whether this full range can be grounded in rules for mechanisms.

22. Van Fraassen famously regarded the rejection of modality as definitive of an empiricist stance toward science. Few have followed van Fraassen. For discussion of van Fraassen's view, see Ladyman (2000) and Monton and van Fraassen (2003).

23. The view that modal beliefs are beliefs about other possible worlds has infiltrated philosophy and distorted the content of modal belief in more complex ways that it would take more time to untangle.

24. The same remarks apply to other forms of modality defined over totalities, see, for example, Loewer and North (forthcoming).

25. One might say here that what I have given is an account of the source of modal belief but not modal fact. In my view, these are the same question.

26. These are the positive and negative sides of what Smolin calls "physics in a box" (2013).

27. See Ismael (2013).

28. Rules for configurations emerge from, and supervene on, those for components. Fix the rules for components and you fix the rules for configurations, but the converse is not true. Fix the rules for configurations, and you fix the rules for components only if we add a combinatorial principle that

guarantees that there is a global possibility for every describable configuration of components.

29. I have been reticent about attributing globalism to proponents of reductive projects because in almost every case the exposition is ambiguous between globalism and the closely related cousin discussed above, in this section. If the distinction is not made, or not made clearly enough, it is very easy to follow Russell's lead and think of the reductive project in globalist terms, handicapping the project beyond recovery. And the very small step from a globalist conception of law to a picture of the metaphysics of fundamental laws that has a wide currency outside philosophy of science has effects that ripple through philosophy.

REFERENCES

Cartwright, Nancy. 1979. "Causal Laws and Effective Strategies." *Nous* 13: 419–37.

Field, Hartry. 2003. "Causation in a Physical World." In *The Oxford Handbook of Metaphysics,* edited by Michael J. Loux and Dean Zimmerman, 435–60. Oxford: Oxford University Press.

Ismael, Jenann. 2011. "Decision and the Open Future." In *The Future of the Philosophy of Time,* edited by Adrian Bardon, 149–68. New York: Routledge.

Ismael, Jenann. 2013. "Causation, Free Will, and Naturalism." In *Scientific Metaphysics,* edited by Don Ross, James Ladyman, and Harold Kincaid, 208–35. Oxford: Oxford University Press.

Joyce, James. 2007. "Are Newcomb's Problems Decisions?" *Synthese* 156: 537–62.

Ladyman, James. 2000. " What's Really Wrong with Constructive Empiricism? Van Fraassen and the Metaphysics of Modality." *British Journal for the Philosophy of Science* 51: 837–56.

Loewer, Barry, and Jill North. Forthcoming. "Probing the Primordial Probability Distribution." In *Time's Arrows and the Probability Structure of the World,* edited by Barry Loewer, Brad Weslake, and Eric Winsberg. Cambridge, Mass.: Harvard University Press.

Maudlin, Tim. 2007. "Causation, Counterfactuals, and the Third Factor." In *The Metaphysics within Physics,* 104–42. Oxford: Oxford University Press.

Monton, Bradley, and Bas C. van Fraassen. 2003. "Constructive Empiricism and Modal Nominalism." *British Journal for the Philosophy of Science* 54: 405–22.

Norton, John. 2007. "Causation as Folk Science." In *Causation, Physics and the Constitution of Reality: Russell's Republic Revisited,* edited by Huw Price and Richard Corry, 11–44. Oxford: Oxford University Press.

Pearl, Judea. 2000. *Causality: Models, Reasoning, and Inference.* Cambridge: Cambridge University Press.

Pearl, Judea. 2001. "Bayesianism and Causality, or, Why I Am Only a Half-Bayesian." In *Foundations of Bayesianism,* edited by D. Corfield and J. Williamson, 19–36. Dordrecht, the Netherlands: Kluwer Academic.

Price, Huw, and Richard Corry, eds. 2007. *Causation, Physics and the Constitution of Reality: Russell's Republic Revisited.* Oxford: Oxford University Press.

Russell, Bertrand. (1913) 1953. "On the Notion of Cause," In *Mysticism and Logic,* 171–96. London: Doubleday.

Sloman, Steven. 2005. *Causal Models: How We Think about the World and Its Alternatives.* New York: Oxford University Press.

Smolin, Lee. 2013. *Time Reborn: From the Crisis in Physics to the Future of the Universe.* New York: Houghton Mifflin Harcourt.

Steward, Helen. 2012. *A Metaphysics for Freedom.* Oxford: Oxford University Press.

Woodward, James. 2003a. *Making Things Happen: A Theory of Causal Explanation.* Oxford: Oxford University Press.

Woodward, James. 2003b. "Critical Notice: Causality by Judea Pearl." *Economics and Philosophy* 19: 321–40.