5

Causation in a Physical World
An Overview of Our Emerging Understanding
Jenann Ismael

5.1 Introduction

There has been an enormous burgeoning of interest in causation across the sciences, and details of various kinds are easy to locate. One can open up a journal in microbiology and be assailed with detailed models of the causal structure of cells and proteins. One can find textbooks on the increasing array of formal and computational tools for causal search and discovery, and take classes devoted to formal methods and techniques in causal modeling. Psychologists are unraveling new details about causal learning, and how people use causal concepts in reasoning. While details of these kinds are in abundant supply, it is not easy for the philosopher or the metaphysician to find among them an answer to the question of what causation is or how it fits into a physics-based ontology.[1]

   This is a survey chapter, or primer, that tries to fill that hole by

 (i) assembling the pieces of our emerging scientific understanding of causation into an account of where (and how) causation arises in the architecture of the cosmos;
(ii) looking at how this emerging scientific understanding illuminates and transforms traditional philosophical questions about causation; and
(iii) considering whether and in what sense the resulting view is pragmatist.

---

[1] One of the charges that is sometimes made against the causal modeling framework, when it is presented as a philosophical account of causation by Woodward, is that the scientific questions somehow sidestep the philosophically important questions. And Steven Sloman prefaces his lovely book *Causal Models*, with an explicit disclaimer that he will not address the metaphysical questions. See Sloman (2005), and Woodward (2017) for Woodward's attempt to address those charges.

5.2 The Back History

If one looks back at the early history of science, it can be seen as developing out of systematization and abstraction of causal thinking: the search for an understanding of the causal relations among events. That is as true of physics as it is of the other sciences. I think that it would not be a mis-description of physics to say that it was devoted, in part, to uncovering the causal skeleton of the world.[2] At first, the notion of cause retained its close connection with mechanical ideas. A cause was something that brought about its effect by a transfer of force. But by the time Newton had finished with it, causation didn't appear in the presentation of his theory at all. What he put in the place of causal relations are mathematical equations that give the rate of change of a quantity over time.[3] Forces drop out of the picture.[4] There is no *direction* or asymmetry in the determination. Later events determine earlier ones as surely as earlier ones determine later ones. In 1913, Russell published a paper – "On the Notion of Cause" (1913) – that is the starting point of almost every discussion of causation and physics. In that paper, Russell noted that the kinds of laws Newtonian Mechanics gives us are so different from causal relations as traditionally conceived that it is misleading to think of them in causal terms at all. And he argued that causation is a folk notion that has no place in mature science. The modern discussion of causation in the philosophy of science really began with Cartwright's deeply influential critique of Russell's paper. Cartwright (1983) 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

---

[2] I take this phrase from Russell, whose views on causation changed between 1913 and *The Analysis of Matter* (1927): "The aim of physics, consciously or unconsciously, has always been to discover what we may call the causal skeleton of the world." For discussion of his later views, see Elizabeth Eames (1989) and Kenneth Blackwell (1989).

[3] Russell uses Newton's law of gravitation as an example: "Certain differential equations can be found, 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. This statement holds throughout physics, and not only in the special case of gravitation." Russell, 1913, first page. For discussion of Newton's treatment of causation, see Janiak (2013), Janiak and Schliesser (2012), and Ducheyne (2011).

[4] Bohm made the same point a little later, in his masterful book *Quantum Theory*: "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 much of 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, 1951, p. 151)

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, i.e., what brings it about. Cartwright's observation led to a lot of handwringing, wondering what causal information adds to the kind of global dynamical laws that Russell took as paradigmatic of physical laws. Philosophers played with probabilistic and counterfactual analyses, and there were a lot of unresolved questions about the metaphysics of causal relations.[5]

Since around 2000, there has been a revolution in our understanding of causal thinking stemming from allied developments in a coalition of fields from computer science to psychology. And it was really the development of the Structural Causal Modeling framework by Judea Pearl (2000) and Spirtes et al. (2000), that transformed the discussion. The framework gives us a precise formal framework for representing causal relationships that is well suited to causal search and discovery in science. It has been at the center of developments in causal learning and statistical inference. It can be used to define normative solutions to causal inference and judgment problems. It has facilitated new insights into the role of causal thinking and cognition. The influence and power of the framework are undisputed and largely due to its utility in raising concrete scientific questions that can be investigated in the laboratory.

From a philosophical point of view, this – in conjunction with developments in physics that I will describe below – has helped us resolve the puzzle brought out by the exchange between Russell and Cartwright: the disappearance of causes from the fundamental level of physical description, on the one hand, and their indispensability in practical reasoning, on the other. The way that it resolves the puzzle is by giving us an articulated understanding of where (and how) causation (or, better, causal thinking and the physical structures that support it) arises in the architecture of the cosmos. And, in the way that is characteristic of philosophical questions once they pass into the hands of the sciences, it has transformed some of the old questions, while also raising new ones.

---

[5] See Hitchcock (2016) and Menzies (2014) for surveys in the development of probabilistic, counterfactual theories. See Field (2003) for a now classic article that set the agenda for metaphysicians trying to understand causation in a physics-based ontology.

5.3 The Formalism

To highlight a few relevant features of the formalism, causal relations are treated, in the first instance, as relations between types of events.[6] Type causation relates types of events like the striking of a match of a specified kind under given circumstances, rather than the particular striking of a particular match at a particular time and place. This accords well with practice in science. A causal relation between types of events is always relativized to a particular network. Networks are defined by collections of variables. Individual properties are represented as the values of variables. The choice of which variables to include and which to leave out make a difference to whether there is a causal link between a pair of variables and what the link is. Singular causal claims (e.g., the claim that a dot on a photographic plate was caused by the impact of an electron) are derived from type causation. DAGs (directed, acyclic graphs) are used to represent the causal relations among the variables in a network. Direct causation, represented by an arrow, is the most basic causal relation in a network. 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. DAGs give us a perspicuous way of displaying causal relationships. Figure 5.1 shows a couple of examples.

You can see at a glance how intervening on one will affect other variables in the network.

From the point of view of the Structural Causal Modeling framework, causal judgments in everyday contexts rely on loose and tacit assumptions about
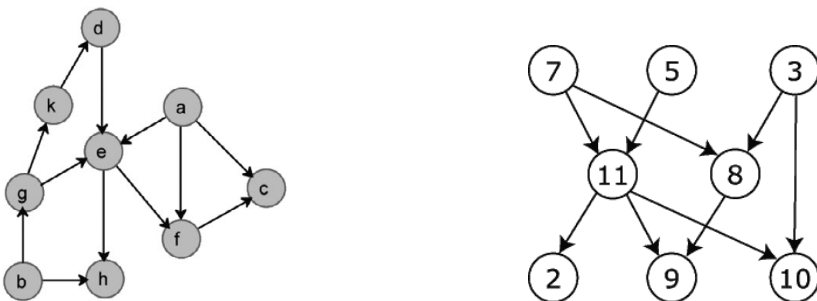

Figure 5.1 Example DAGs.

[6] See Pearl (2000). Also Pearl (2009) and Woodward (2003).

which variables are being considered, and what is being held fixed in situ for purposes of assessing the effects of one variable on another. Those choices are made explicit in the formalization.[7] Because it makes them explicit, it lets us separate the objective content of causal judgments from the context-dependent pragmatic factors that guide the choice of network. This means that once we have specified the choice of network and contextual factors, it is a factual question whether there is a causal relation that exists between a pair of variables, but without that specification we don't have a well-defined question of fact.

The question that comes out of the Russell/Cartwright exchange is:

How does causal information go beyond the information captured in dynamical laws, i.e., what does causal information add to those laws, or (if it doesn't add anything) what information contained in those laws does it make explicit?

Global dynamical laws give us information about relations between global states at different times. Causal information is information about the results of hypothetical interventions, where by that we mean information about what would happen if a variable is separated from its past causes and allowed to vary freely. The reason that this is not information that we can in general extract from global laws is not that global laws *don't* have modal content, i.e., it is not that they don't go beyond making claims about what is actual and have implications for what is possible. The problem is that the modal content that they have is global. Global laws tells us what kinds of world are possible. They don't tell us what would happen if we intervened in one of those worlds in a way that is contrary to law. (In logical terms, this is put by saying that the antecedents to intervention counterfactuals are nomologically impossible, i.e., they violate the physical laws.)

That gives rise to another question: how is causal knowledge possible?

If causal information is information about what would happen if some variable were separated from its causes and allowed to vary freely and no variable is ever actually separated from its causes, how could we study causal structure? The answer is that, in practical terms, there are various ways of randomizing the values of a variable, which provides an effective separation from its own value, so that we can treat changes to the values of some variables as interventions; e.g., insert a coin flip, have a random number generator set the value of an exogenous variable, or let a lab assistant choose them at will. Any variable that can be manipulated and whose own causes are uncorrelated with the variables of interest can be regarded, for such purposes, as external to the

---

[7] For these reasons, the Structural Causal Modeling framework is not an analysis of the folk concept of cause.

system, and changes in the value of such variables as being interventions.[8] This is the essence of scientific experimentation. The possibility of knowledge rests on the fact that we can effectively isolate systems in the laboratory, manipulate the external input to them, and observe the effects of our manipulations.

And this in its turns tells us something important about how we should think of (for lack of a better word) the metaphysics of causation (or better, the physical structure that support causal beliefs, i.e., the physical structures that causal beliefs carry information about).

In a world that is as complicated as ours, and which decomposes into a collection of rule-governed components, which can be isolated and studied in the lab, and where the laws governing configurations are given in terms of the laws governing components and their interactions, causal information can be extracted from laws governing components and interactions. One can concoct examples of very simple worlds, or worlds in which there are nomological restrictions on initial conditions, in which causal information can*not* be extracted from *global* laws. What is really going on in those worlds is that the global dynamical laws will not allow us to recover the laws governing their components and interactions. So long as one is willing to take the local nomological structure of the world as basic, one can ground facts about the results of hypothetical interventions in that (and this follows standard practice in physics: it's a perversion of the content of physical theories to make a lunge for the global laws, and then insist that the full modal content of Newtonian Mechanics has to be somehow extractable from the form they take in our world).

5.4 Surface Structure and Underlying Processes

The relationships captured in DAGs ("interventionist causal relationships") are experimentally accessible, and often first in the order of discovery. They are part of the surface fabric of the world, and almost tailor fit to play the role that Cartwright highlighted when she argued for the indispensability of causal notions for practical reasoning.[9] But there is another tradition, which comes out of the later Russell (by way of Reichenbach and Salmon) and receives

---

[8] 'Freedom', for our purposes, can be taken to mean simply 'uncorrelated with variables of interest.'

[9] In fact, it is very natural to use that practical role as an implicit definition for causal relationships in analogy with the role that the Principal Principle plays for chance. See Lewis (1980), reprinted with added postscripts in Lewis (1986) and Price (2012a). On this understanding, causal relationships are those relationships that serve as strategic routes for bringing about ends.

its fullest development in Phil Dowe's work (see Dowe (2008)). This tradition analyzes causation in terms of causal processes. Causal processes are chains of events related by causal *interactions*, where a causal interaction is an interaction that involves the local exchange of a conserved quantity. Causal processes can be defined more directly in terms of micro-dynamical laws. Causal process and interventionist accounts are sometimes treated as competing accounts of causation, but we don't need to choose between them. They are both useful, and in classical physics we find examples of both. We have causal processes at the fundamental level, and manipulable causal relationships of the kind captured by DAGs all over the place. These are undiscriminating about levels, and neutral about underlying processes. They represent often local, scaffolded relationships among variables that can be investigated experimentally, frequently without knowledge of the underlying processes.

Both, moreover, capture some aspects of everyday notions.[10] The interventionist notion captures the idea that you have a causal connection when one thing can be used to bring another about. The causal processes capture the idea of influences that propagate through space, and they satisfy the expectation that causal connections should be mediated by mechanisms, and involve some sort of contact. If interventionist causal relationships are part of the surface fabric of the world, causal processes are part of the deep, theoretically postulated substructure that supports them.

5.5 Temporal Asymmetry

I stated that classical physics gives us both causal processes and the high-level causal networks captured in DAGs, but there is a puzzle about the relationship between the sorts of high-level causal pathways and the underlying dynamical laws that is crucial to understanding causal notions of both kinds.

Interventionist causal relationships of the kind captured in DAGs have one striking feature: they are temporally asymmetric. They run in one direction – from past to future. The temporal asymmetry is no mere matter of convention. If we follow Cartwright in saying that what is definitive of causal relationships (as distinct from mere correlations) is that they can function as strategic routes bringing about ends, it is a manifest and pervasive empirical fact about our world that all known strategic routes to bringing about ends run from past to future. Of all the reasons that Russell gave in that early paper for saying that causation had disappeared from physics at the fundamental level, this is the

---

[10] Neither captures every part of the everyday notion, which is a mix of elements drawn from causal experience, proto-theories, etc.

one that persuaded most people. The underlying dynamical laws are simply equations of motion that relate the state of the world at one time to its state at others. They have no intrinsic direction of determination.

Now, this gives rise to a straightforward physical puzzle: How do the kinds of high-level causal pathways that are an objective part of the surface fabric of our world arise from temporally symmetric underlying laws? This is not the kind of question that people in the lab studying signaling pathways in cells typically ask. Scientists who empirically investigate causal structure take it for granted that causal relations run from past to future. They set up experiments that allow localized interventions and look for the temporally downstream effects.

But it *is* a theoretical question for the physicist trying to relate the kinds of causal pathways studied in biology labs to an underlying theory. And as it happens, it has been illuminated by research into the foundations of thermodynamics. Such research began in the mid-nineteenth century and was originally focused on trying to derive the phenomenological asymmetries embodied in the second law of thermodynamics from the time-symmetric laws of classical mechanics. It only quite recently turned into a very general account of the sources of temporal asymmetry in our world. So a conversation that started about why gas will disperse to fill an open container but not collect spontaneously in one corner, and why ice melts when placed in a glass of warm water but we don't see it spontaneously form without a change in temperature, turned into a conversation about why we remember the past but not the future, and why time seems to flow from past to future, and why causation doesn't run backwards. And – as in the revolution in our thinking about causation that came with the development of the Structural Causal Modeling framework – the input from both scientists and philosophers was instrumental. There was Boltzmann and Maxwell and Gibbs (all scientists), but also Reichenbach (1956), Huw Price (1996) and David Albert (2000) (all, of course, philosophers). These are two places where the interaction between science and philosophy has been a model of productive inquiry.

There were an enormous number of false steps, and although there are many unanswered questions, David Albert's 2000 book (*Time and Chance*) brought a conceptual focus that made it possible to draw philosophical conclusions. Let me give you a very quick sketch of what Albert's account looks like.

It has three first principles:

1. the Newtonian laws of motion...
2. a statistical postulate
3. something he calls the Past Hypothesis.

The Newtonian laws of motion are the familiar time-symmetric laws that Russell was referring to. The statistical postulate is a probability distribution that assigns probability to a system's being in a given microstate, given its macrostate, that is a strict addition to the laws (it doesn't follow from them directly, though there are a lot of attempts to motivate it dynamically). The Past Hypothesis is a contingent hypothesis about the early history of the world. It says that the universe started in a state of low enough entropy to make thermodynamic generalizations applicable for the roughly 15 billion years we think these generalizations have held. The only thing you have to know about entropy, for immediate purposes, is that the statement that entropy increases but never decreases embodies all of the temporally asymmetric phenomena that fall under the scope of thermodynamic description, and maybe all of the known temporal asymmetries in our world.

The three postulates work together as follows. The Newtonian laws delimit the space of physically possible worlds. The Past Hypothesis knocks out all of those worlds that don't start in a low entropy state. If you take the initial microstates of those worlds and apply the Newtonian dynamical laws, you will see that some of the remaining worlds will be on entropy increasing trajectories, and some will be on entropy decreasing trajectories. The statistical postulate provides a probability distribution over remaining worlds that overwhelmingly favors those on an entropy-increasing trajectory. The result is that the most probable history of the universe is one wherein entropy rises.

There are a couple of reservations that one might have about this story.

Albert tells this whole story at the global level. But one need not do that. Formally, it works just as well for local adiabatically isolated subsystems of the world (systems that aren't exchanging energy with the environment), and there are reasons for thinking it might be best told at that level: (i) those are the systems to which thermodynamics is applied, (ii) it is not obvious whether the local story follows from the global one, and (iii) the probabilities have a clearer interpretation. But this issue doesn't matter here.

In addition, if you look at Albert's postulates, the Newtonian laws are time symmetric, and the statistical postulate is time independent, so the only time-asymmetric assumption is the Past Hypothesis. It's a particularly crucial part of the story because it's doing all of the explanatory work. There has been some important dispute about the right form for this hypothesis. Everyone agrees that some kind of temporal asymmetry in boundary conditions needs to be assumed to extract the temporal asymmetries of thermodynamics. In Albert's account, it is a low entropy state for the universe at some time in the distant past. Mathias

Frisch (2014), in a discussion that is specifically devoted to the asymmetries embodied in causal reasoning, has shown that there are different ways of expressing the time-asymmetric assumptions that are needed to generate the sorts of asymmetries that would support causal inferences (I think that it is right to posit that this is really the root asymmetry and that if you had it in place, the asymmetries of time and information would follow). Frisch points out that it would do just as well, for example, to assume a causal Markov condition (that any node in a DAG is conditionally independent of its non-descendants, given its parents), or an assumption of microscopic randomness in the initial, but not final state of a system. There is no substantive question as to which of these assumptions is more fundamental. They are different ways of describing the structures and the contingent asymmetries that in our world support causal inferences running from past to future. David Wallace (2012) has also argued (with a more foundational focus, but in a way that ends up making points allied to Frisch's) that in the explanatory story, what is needed to generate thermodynamic generalizations is not really the entropy but the randomness or uniformity of the early universe. So while there are still things to sort out, there is broad agreement that (pending the mathematical conjectures on which the account rests) the local macroscopic asymmetries in our part of the universe could have their origin in the state of our universe in the distant past, which developed under (effectively) classical laws into a world with the emergent macroscopic asymmetries captured in thermodynamic generalizations.

If one is interested in how this story connects to causal inference, one will be interested in extracting a Causal Markov condition.

But if one is interested in how this story connects to a notion of causation implicitly defined by Cartwright's observation that if you're given a body of correlations, the causal pathways are those that can be used as strategic routes to bringing about ends, here's how to get something that connects more naturally to that idea.[11] If we start with a thermodynamic gradient, and we hold fixed all of the information embodied in the present surveyable macroscopic state of the world (and here, I don't mean just looking around you and holding fixed the ambient temperature and dispositions of, for instance, medium-sized dry goods, I mean holding fixed all of the information in libraries and scientific databases: all of the information that we've somehow accessed and recorded, or

---

[11] This is a conception of causation that brings a connection to agency into the foreground and connects to people's pre-theoretic ideas about causation. I don't think there is a fruitful discussion to be had about what the one true conception of causation is. It's part of the general approach here to show how these variously articulated concepts arise, exhibit the practices in which they figure, and the external and internal structures that support those practices.

that we could have accessed and recorded), and we ask about the probabilistic effect of local interventions on its present state (of a kind that are possible in practical terms for agents like us), it turns out that they will propagate asymmetrically into the past and future. They will, that is to say, make a difference to the future, but leave the macroscopic past untouched. And what this means is that if we look for the strategic routes for bringing about ends – i.e. if we look for ways of doing something in the here and now that will make a difference to the probabilities of events at other times and places – we will find that there's nothing one can do to make a difference to the past, but a lot we can do to make a difference to the future (ignoring, of course, Cambridge differences such as making a liar of yourself, or redeeming early promise from later missteps. . .)

5.6 Causal Concepts, Causal Learning, Causal Experience

So far, everything that I have said stays at the level of physical structure. We have talked about the temporally symmetric classical laws, and the emergent asymmetric interventionist pathways that can be studied in the lab and captured in DAGs. So, we have the external infrastructure in place that describes the macroscopic fabric of the world. Now, we are in a position to talk about causal *concepts*. We've moved from talking about structure in the world to talking about something that has a role in an internal cognitive economy. To do that, we introduce an agent and we couple it to the macroscopic structure. So, we think of it as a system whose job is to extract information from the environment and use it to guide behavior, and we specify that the information it takes in is macroscopic. And we look at how causal concepts arise, which is to say that we look at the role they play in an internal inferential network that is linked to experience at one end and action at the other. The sciences that are relevant are psychology or cognitive science, and there is a wealth of fascinating work on causal cognition and learning, on the development of causal ideas, their role in inference, decision and action, as well as moral judgment.

If we step back and take a cross-section of the world, ordered by scale, and look at the layers of structure from the microscopic all the way up to the level of the human being interacting with a macroscopic environment, the picture that comes into focus appears something like this: There is the geometry of space-time, which (in a classical setting) imposes (or embodies) constraints on the causal connectibility of events. There is the matter content, the temporally symmetric microscopic laws that tell us how the state of the world at one time relates to its state at the next, and the emergent macroscopic asymmetries.

The macroscopic asymmetries support the emergence of creatures that use information to guide behavior. Eventually, they get complex enough that they do something that we would recognize as thinking. By the time we get to creatures like us, thinking is causal from stem to stern.[12]

5.7 Assessment

While there is much that is not known, and various points of dispute, the sketch that I have given here is robust with respect to most of the emerging details.[13] Notice how much it transforms both our understanding of causation and of time. It is very different from anything that might have come out of simply reflecting on the concepts. No single part of this story is the story of 'what causes are.' But I think that once all of these pieces are assembled, there is no further question that remains to be answered, no story about the metaphysics of causation that remains to be told.

There are two questions that philosophers always want to ask about causes:

(i) Where does the asymmetry 'come from'?
(ii) Are causes inside or outside the head; i.e., are they part of the mind-independent fabric of reality, or in the eyes of the observer?

We've just seen a complex answer to the source of the asymmetry. What we observed was that there is not a single asymmetry, but rather a couple of related asymmetries. The thermodynamic gradient created by a low-entropy state of the early universe creates a *macroscopic* asymmetry that supports the emergence of information gathering and utilizing systems. When such systems get sufficiently complex, information gathering and utilization takes the form of conceptual activity. In this way, creatures develop causal concepts with a built-in asymmetry that reflects an asymmetry in their practical and epistemic relationships to the world. So while there is no fundamental asymmetric relation of determination in physics,[14] for creatures like us, who couple to the macroscopic structure of the world, the fact that information flows from the past, and causal influence runs into the future is a fundamental and formative

---

[12] What I have described is the part of the stratigraphic hierarchy that is relatively well understood in physical terms. If we look below the Planck scale, the theory that we look to is quantum mechanics, and things get confusing and too unsettled to say anything definite.

[13] See Frisch (2014). The sketch is robust with respect to many of the details; it does cede to any new scientific developments.

[14] Questions about temporal asymmetry in both classical and quantum mechanics, are subtle and there is room for contention, but not in ways that make a difference to our discussion. See Albert (2000), North (2008), and Arntzenius and Greaves (2009) on temporal symmetry in classical mechanics, and Bacciagaluppi (2006) on the quantum case.

feature of our experience, one that frames every aspect of our thoughtful engagement with the world.[15]

On the question of whether causes are inside or outside the head, the right thing to say is 'partly inside, partly outside.' Our pre-scientific ideas about causation aren't articulated enough to single out one part of layered structure that leads from microscopic processes up to the coupled interaction between agent and environment in which causal thinking has its home as 'where the causes are.' Scientifically, there is no uniform usage; physicists will speak of the spatiotemporal geometry as encoding causal structure, but when biologists talk about causal structure they mean something much higher up in the stratigraphic hierarchy. When ordinary people talk about causation, they invoke ideas closely associated with the phenomenology of pushing and pulling, something with a quasi-muscular notion of compulsion in mind. We can observe how all of these ideas arise and relate to one another. I don't see that there is a factual question to be settled about which of them captures the essence of causation, and at what point the real causes come in.

Notice that this kind of account is different from what we might think of as the prototypical philosophical analysis. It's been given different names. Huw Price (2011) calls it a "genealogy," Doug Kutach (2014) calls it an "empirical analysis." It is characteristic of the way in which we get a scientific account of how these notions play a central role in mediating our practical and epistemic interaction with the world. There's an external component (the infrastructure in the environment that agents use as strategic pathways to bring about ends) and an internal component (an account of the content of causal concepts understood partly in terms of their connection to experience and action, and partly in terms of their role in an internal inferential network). A big part of fitting those two pieces together is saying how agents with the kind of information we have, and with our practical capacities for intervention, in a world that is objectively structured like ours, develop and use causal ideas.

There are two philosophically important aspects of the account: (i) it secures the objectivity of causal relations, making them apt objects of scientific study, and (ii) it's a good antidote to reifying causal relations as compulsive asymmetric relationships between localized events. There are causal pathways written into the fabric of the world, but those pathways are neither asymmetric

[15] Although the body is a physical system, and interacts microscopically with the environment, the mind is an information-processing system that receives information through perceptual channels. To say that we 'couple to the macroscopic structure of the world' I mean that the mind receives information about the current macroscopic state of the environment, and uses that information to guide interventions. From the mind's perspective, its interventions are macroscopic, because it only gets perceptual feedback on the macroscopic movements of the body, and can only *control* (in the relevant sense) what it can perceive.

nor compulsive. We see the complete account of the metaphysics of causation as given by the facts about the world and our position in it that support the acquisition and use of causal concepts. Furthermore, this account does NOT treat causal relationships as asymmetric, pushy relations of dependence that are written into the fundamental fabric of the world. But it makes causal relationships objective in all of the ways that should matter to science. It makes perfect sense to think of science as devoted to (as Russell put it in his later years) discovering the causal skeleton of the world.[16]

5.8 Conclusion

The question of how causation arises in the structure of the cosmos has been a live one since Russell observed that it appears to have disappeared from the fundamental level of physical description. In this chapter, I sketched an answer to that question that I take to be emerging from what has been learned about causation from allied developments in a coalition of fields from computer science and psychology to physics. Questions about causation span disciplines, and the workshop that produced this volume was a wonderful illustration of the need for interdisciplinary thinking in addressing them.

Acknowledgments

---

[16] There's a particularly interesting series of exchanges between Woodward and Price, who agree on the contours that I have sketched, on whether, and in what sense, causation is objective or 'perspectival' (Price's word for agent-dependent) on a broadly interventionist account of the content of causal claims. Woodward (2013, 2009) is concerned to protect the objectivity of causal claims, and Price (2005, 2017) points out the ways in which causal direction is relative to the agent's capacities for intervention. See also Ismael (2016).