1.4 The Causal Element in Explanation

1.41 The Direction of Explanation and the Direction of Causation

Certain well-known counterexamples to the DN account provide strong reasons to prefer a causal approach to explanation over either an expectability or a pattern subsumption approach.

Given the height of a flagpole, the position of the sun, and a simple physical law—light travels in straight lines—you can deduce the length of the shadow cast by the flagpole. Thus, according to the DN account, you can explain the length of the shadow by citing the flagpole's height, the position of the sun, and the rectilinear propagation of light.

In the same way, however, you can deduce the height of the flagpole from the law, the position of the sun, and the length of the shadow. Thus, according to the DN account, you can explain the height of the flagpole by citing, among other things, the length of its shadow. Contrary to the DN account, however, we do not consider the length of the shadow to explain the height of the flagpole. Therefore, the DN account is mistaken.[8]

How ought the DN account's troubles to be diagnosed? There is an asymmetry in the two explanations where there is no asymmetry in the arguments. This suggests there is something more to explanation than mere argumentative force, and that this something is a relation that is asymmetric in the case at hand—a relation that holds between flagpole and shadow, but not between shadow and flagpole. The obvious candidate is the relation of causation. The flagpole causes the shadow to exist but the shadow does not cause the flagpole to exist. Explanation, it seems, flows in the same direction as causation.

This conclusion is reinforced by another famous counterexample to the DN account. Suppose that, whenever the needle on a barometer dips, there is a storm. From this meteorological law and a particular needle dip, you

8. The argument is due to Sylvain Bromberger; however, the now-universal choice of flagpole as paradigmatic shadow-caster seems to have been a later innovation (Salmon 1990b, 189).

can deduce that a storm will occur. Thus, according to the DN account, you can explain the storm by citing the law and the fact that the needle dipped. Contrary to the DN account, however, we do not consider the barometer reading to explain the ensuing storm. Therefore, the DN account is mistaken.

Again, it seems that the existence of a sound deductive argument is not sufficient for explanation. What is also required is, as in the case of the flagpole, a causal relation between the explanatory facts and the explanandum. What the barometer does not cause, it cannot explain.

1.42 Towards a Causal Basis for Explanatory Asymmetry

The counterexamples discussed in section 1.41 provide good reason to think that the explanatory relation has a causal element (for two alternative views, see section 1.47). Let me outline what I take to be more or less the minimal assumptions about causality necessary to handle explanatory asymmetry.

Take as a starting point a schematic causal constraint on explanation, which counts as the explainers of an event *e* only entities satisfying one of the following three criteria:

1. An explainer can be an event or state of affairs that bears a certain causal relation to *e*. Without prejudice, call this the relation of *causal influence*. Thus, to qualify under this first criterion, a potential explainer should be a causal influence on *e*.

2. An explainer can be a causal law in virtue of which an event satisfies criterion (1), that is, a causal law in virtue of which an event causally influences *e*.

3. An explainer can be a background condition in virtue of which a causal law satisfies criterion (2), that is, a background condition, satisfaction of which is necessary for an event to causally influence *e* in virtue of a certain causal law.

Call any event, law, or set of background conditions that satisfies one of the criteria a causal influence on *e*. In this sense, the causal influences include not only the events connected to *e* by the causal influence relation, but the laws and background conditions in virtue of which such connections exist—a slight abuse of terminology, perhaps, but economical. You can then say that all explainers of *e* must be causal influences on *e*. In what follows, I will focus on potential explainers that qualify under clause (1).

In the above scheme, *causal influence* is a mere placeholder: it is an open question what kind of thing the influence relation might be. If the causal constraint on explanation is to deal with the flagpole/shadow problem and other similar cases, however, the influence relation take a certain form. In particular, the influence relation must have the properties of *asymmetry* and *particularity*.

An influence relation is *asymmetric* if it is possible for causal influence to run in one direction only, that is, for *c* to causally influence *e* without *e*'s influencing *c*. (Mutual influence is possible, then, but not necessary.) It is the asymmetry of influence that allows us to distinguish, on causal grounds, between the nomological dependence of the length of the shadow on the height of the flagpole and the nomological dependence of the height of the flagpole on the length of the shadow, declaring one but not the other to be a relation of causal dependence, and so a potentially explanatory relation. Typically (though perhaps not always), the causal influence relation should run from past to future, providing the undergirding for the notion that states earlier in time causally explain, but are not explained by, states later in time.

Asymmetries in causal influence should be found everywhere, since in almost every process—evolutionary, ecological, economic—science distinguishes a causal order and holds that explanation must not run counter to this order. It is perhaps too much to say, yet, that every explanation must follow the direction of causal influence, but certainly, in every domain there are explanations that must do so. This is true even of fundamental physics, since the intuitions of explanatory asymmetry summoned up by flagpole/shadow cases do not disappear at the fundamental level. The positions of Rutherford's scintillations—showing the final positions of alpha particles passing through gold foil—no more explain the structure of the gold atom than the shadow explains the dimensions of the flagpole. But the quantum scattering processes involved are just about as fundamental as you can get. Analogs to the barometer/storm case are equally easy to find. (Skepticism about the applicability of causal concepts to fundamental physics will be discussed in section 1.5.)

So much for asymmetry. An influence relation between events is *particular* to the degree to which it relates very finely individuated aspects of the world. An influence relation that is not particular at all might simply relate the complete state of the world at one time to its complete state at a later time (Maudlin 2007). Such a relation provides enough discriminatory power to deal with flagpole/shadow cases, because the length of the shadow at any moment depends on the height of the flagpole at a slightly earlier time, and so an attempt to explain the former in terms of the latter goes contrary to even this very coarse-grained causal relation. But it is less well able to deal with barometer/storm cases, since here the spurious explainer, the barometer reading, temporally precedes the event to be explained; it therefore belongs to a state of the world that causally influences the state of the world to which the explanandum belongs, so its explanatory status cannot, if influence is unparticular, be called into doubt on purely causal grounds.

A causal influence relation that is very particular runs, by contrast, between very fine-grained aspects of earlier and later states of the world; it therefore allows you to discriminate between those aspects of a world state that do and those that do not causally influence a later storm. (That said, the presence or absence of causal influence alone may not be enough to identify all barometer-like spurious explainers; see section 1.43.) Because the causal influence relations found in the literature are, Maudlin's excepted, all very particular, I will say no more about particularity in later sections; it will be there in the background but unacknowledged.

What next? A natural strategy is to find the best available account of causation that provides a causal influence relation satisfying the requirements of asymmetry and particularity. To study explanation, then—or causal explanation at least—is first to study the metaphysics of causation. Causal accounts of explanation in the literature have tended to take just this approach, running up their flags over some or other metaphysical theory of causation before exploring the terra incognita of explanation (Salmon 1984; Lewis 1986a; Woodward 2003).

I will do something different. There are several attractive accounts of causation that provide a relation of causal influence that is both asymmetric and particular. Why choose? I propose a provisional causal ecumenism. As the investigation proceeds, this ecumenism will go from provisional to provincial, but none of the accounts of causation I am about to describe will be declared entirely incompatible with the kairetic account of explanation.

Three views of causation well equipped to supply explanation with its particular asymmetry are Dowe's conserved quantity account, Lewis's various counterfactual accounts, and Woodward's manipulation view. Let me say something about each, focusing, for now, on the deterministic case. (I should note that, when I say that these accounts are well equipped to handle the explanatory asymmetry problem, I do not mean that they are free of problems, let alone that they are fundamentally correct in their approach to causation. I mean that the various kinds of causal relation they posit, *if they are correct*—and presumably at most one can be correct—are adequate to deal with the asymmetry problem.)

1.43 The Conserved Quantity Theory

Dowe (2000), building on the work of Reichenbach (1956) and Salmon (1984), defines an asymmetric relation of *causal connection* based on the notion of a conserved quantity, that is, a quantity preserved by the physical laws, such as mass-energy or charge. Two objects-at-times are causally connected just in case a line can be drawn in space-time from one to the other that at every point either (a) follows the world line of a single object, or (b) switches from the world line of one object to the world line of another at a point where the lines intersect and where the two objects exchange a conserved quantity. Causal connection, then, is by any combination of *persistence* and *interaction*, where persistence is simply the relation between an object at one time and the same object at any other time, and interaction is the exchange of a conserved quantity by two objects in the same place at the same time.

An electron at time t_1 might be causally connected to an electron at time t_2, for example, because at some time after t_1 the first electron emits a photon (the bearer of electromagnetic force) that influences the second electron before time t_2. The causal connection goes from the first electron at t_1 to the same electron at the time of the emission by persistence, from the electron to the photon at the time of the emission by interaction, from the photon at that time to the time when the photon interacts with the second electron by persistence, and so on. This, according to Dowe, is the reason that what appears to us to be a repelling of one electron by the other (and vice-versa, of course) is indeed a genuine causal connection.

The direction of causation, usually but not always forwards in time, is not determined by the nature of either persistence or interaction, but by a separate element of the account based on one of Reichenbach's ideas, which I will not describe here (Dowe 2000, chap. 8). Related accounts of causation have been offered by Fair (1979) and Salmon (1984); as a group these go under the name of *process theories* of causation.

Dowe's account offers a definition of a symmetric causal connection relation between objects as well as a definition of the local direction of causation, but to solve the explanatory asymmetry problem, what is wanted is an asymmetric causal influence relation between events or states of affairs. No problem; count an event c as causally influencing an event e just in case c and e involve objects that are causally connected and, if they occur at different times, the connection from *c* to *e* at all times goes in the local direction of causation.

This gets you what you want with the flagpole and shadow: the direction of causal influence is never from the shadow to the flagpole. The case of the barometer is slightly more complicated. It is tempting to say that there is no causal influence between the barometer and the storm, ruling out the explanation of one by the other. But this is not quite correct: the parts of the barometer exert a minute gravitational influence on the parts of atmospheric whole that is the storm. The simple reaction is to disregard the influence on the grounds that it is so small as to be negligible. To take this course, some way of quantifying causal influence is required. You might, for example, set the degree of influence equal to the amount of conserved quantity exchanged, relative to the whole. For reasons that will emerge in chapter two, I think that this simple reaction is not the last word on the explanatory significance of minute causal influences, but let it stand for now.

Note that on Dowe's account, the causal connection between any two objects or events, such as barometers and storms, will ultimately depend on the causal connections between their fundamental constituents—the fundamental particles that either make up the object, or whose behavior realizes the event. In particular, two events are causally related just in case any two fundamental constituents of the events are causally related (and the direction of causation is right). In this sense, Dowe's account can be said to be a *fundamental level* account of causation: all causal relations go by way of the fundamental level. The other two views of causation considered in this chapter will provide a contrast.

1.44 The Counterfactual Theory

According to the simple counterfactual theory of causation, an event *c* causes an event *e* just in case, had *c* not occurred, *e* would not have occurred. An icy patch on the road is a cause of a car crash, then, because if the icy patch had not been there, the crash would not have happened. But it is not true that if the crash had not happened, the icy patch would not have been there. So the icy patch causes the crash, but the crash does not cause the presence of the icy patch. When the effect is an ongoing state of affairs, such as a shadow's being a certain length, you have to be more careful of your evaluation of the counterfactuals; in particular you have to be careful not to use "backtracking" counterfactuals (Field 2003, 448–450). When precautions are taken, you get the asymmetric relation you want: the shadow counterfactually depends on the flagpole, but not vice-versa.

Lewis (1973a) offers an account of causation based on the simple counterfactual theory: *c* causes *e* just in case there is a chain of counterfactual dependence relations connecting *c* and *e*. For my purposes in this chapter, there is no need to distinguish Lewis's from the simple account; for more on the difference, see section 2.4.

The counterfactual account has what is in one sense a far more selective conception of the causal relation than Dowe's account.[9] Events that are causally connected in Dowe's sense are often not causally connected in the counterfactual sense. Suppose I let out a whoop while you hurl a cannonball at a window. The window breaks. On Dowe's view, my whoop and your cannonball-hurling are both causally connected to the window's breaking— my whoop because the sound waves transfer energy to the window as it breaks, by causing it to vibrate. But had I not whooped, the window would still have broken, so the whoop is not a cause on the counterfactual account.

Nevertheless, the counterfactual account can make a certain kind of sense of Dowe's claim that there is a causal connection between the whooping and the breaking, as follows. Consider the concrete events that realize the events of the whooping and the breaking. A concrete event, recall from section 1.22, is an event individuated by all its intrinsic properties. Thus the concrete realizer

9. It is, in another sense, far less selective: as explained in section 6.32, it counts vast numbers of omissions as causes.

of the whooping is the event of the whooping's happening in exactly the way it did, down to the last ululatory waver, and the concrete realizer of the breaking is the event of the breaking's happening in exactly the way it did, down the shape, size, and trajectory of the smallest shards of broken glass. The realizer of the breaking counterfactually depends on the realizer of the whooping: if I had not whooped, the breaking would not have happened in exactly the way it did, since the sound waves I generated did have some slight influence on the movement of the glass molecules. Say, then, that one event *causally influences* another if, had the concrete realizer of the first not occurred, the concrete realizer of the second would not have occurred.

This causal influence relation is not, note, a different kind of relation from the counterfactually defined *is a cause of* relations between high level events: it is simply the subset of the *is a cause of* relation that has as its relata only concrete events. I have introduced a new name for some of the *is a cause of* relations admitted by the counterfactual account, rather than adding anything new to that account. Now, a charitable proponent of the counterfactual account of causation can interpret Dowe's claims as being about counterfactual dependences between concrete events, that is, as being about causal influence relations.[10]

I called Dowe's account of causal connection a fundamental level account. The counterfactual account can be said to contain a fundamental level account of a certain kind of causal connection—the influence relation—but much more besides, namely, the account of *is a cause of* relations between high level events. For this reason, it might be called a *multilevel* account of causation, an account that posits distinct causal relations at all levels. Although the causal

10. Recently, Lewis has defined a relation of causal influence much like the relation characterized here (the main difference being, as you might expect, that influence is a matter of chains of dependence), and has tried to give an account of causation between high level events in terms of influence—in effect relinquishing the multilevel approach to causation and opting for a fundamental level approach after all. See Lewis (2000), and for arguments against the account, Strevens (2003a) among others.

relations at one level—between biological events, say—will not be entirely independent of the causal influence relation at the fundamental level, since the fundamental laws of nature have a hand in determining both, the higher level causal relations are not in any simple way made up of, let alone identical to, the influence relations.

Now to my main point: to solve the problem of explanatory asymmetry, only the relation of causal influence is required, or in other words, you only need look at *is a cause of* relations between concrete events. You could handle the asymmetry problem by looking at higher level *is a cause of* relations, but you do not need to: that part of the counterfactual account that overlaps substantially with the Dowe account is sufficient for the job. Either Dowe's or the counterfactual theorist's version of causal influence will asymmetrically relate flagpoles to their shadows, cold fronts to storms, but not shadows to flagpoles or barometers to storms.[11]

1.45 The Manipulation Theory

Woodward (2003) proposes that causal facts are equivalent to, or have as their basis, facts about in-principle manipulability (an objectivist revision of the more anthropocentric manipulability theories of Gasking (1955) and Menzies and Price (1993)).[12] In the first instance, Woodwardian causal relations connect quantities, such as the height of a flagpole or the length of a shadow. To say that flagpole height causes shadow length is to say that the length of a shadow can be changed by manipulating the height of the corresponding flagpole. To say

11. For a strictly correct treatment of the barometer case, the counterfactual theorist's notion of causal influence must, like Dowe's, admit of degrees, so that the influence of the barometer on the storm can be declared negligible. It is appropriate here to look to Lewis's own recent attempt to quantify influence (see note 10), which turns on the degree of change in the concrete realizer of the putative effect *e* that would be obtained by making small changes in the concrete realizer of the putative cause *c*. Similar remarks apply to Woodward's views below.

12. Woodward has expressed reservations about the treatment of his theory as an account of the nature of causality; many philosophers regard a manipulability metaphysics of causation as plausible and promising, however, and it is in this spirit that I invoke Woodward's name.

that shadow length causes flagpole height is to say that the height of a flagpole can be changed by manipulating the length of its shadow. The first claim is true; the second false. Thus flagpole height causes shadow length but not vice-versa. (What constitutes a manipulation is carefully defined by Woodward; the details do not matter here.)

A relation of singular causation between events is derived from the type relation between variables. The form of the derivation will be examined in somewhat greater detail in section 2.5; for now a counterfactual précis, suggested by Woodward himself, will suffice. The manipulation account declares *c* to be a cause of *e* just in case the following counterfactual is true *for some putative causal path between c and e*: if *because of an intervention c* had not occurred, *and had all variables not on the putative path been held constant*, *e* would not have occurred. It is in virtue of the italicized phrases that Woodward's counterfactuals differ from those invoked by Lewisian approaches to causation.

The manipulation account of causation is a multilevel account, in the sense defined in the previous section: it posits distinct causal relations at all levels, independently relating the variables of fundamental physics, biological variables, economic variables, and so on. I want to show that the causal relations the manipulation account posits between fundamental level variables are sufficient on their own to solve the problem of explanatory asymmetry. As with Lewisian approaches, the key is to focus on the causal relations between the concrete realizers of high level events.

Say that a high level event *c* causally influences another high level event *e* just in case the concrete realizer of one is a cause of the concrete realizer of the other, which is to say, just in case there is a causal pathway from one realizer to the other with respect to which the Woodwardian counterfactual holds—just in case, had a manipulation prevented the occurrence of *c*'s realizer, and had all variables not on the pathway in question been held to their actual values, *e*'s realizer would not have occurred. So defined, the manipulationist's causal influence relations are of a piece with the other causal relations between events posited by the manipulation account; they simply constitute a subset of the full complement of manipulationist singular causal relations. This subset is enough, I claim, to solve the explanatory asymmetry problem: the influence relation will hold between flagpole and shadow but not vice-versa, and so on. The close correspondence between Lewis and Woodward counterfactuals should be enough to convince you that this claim is at least roughly correct.

1.46 Causal Ecumenism

The conserved quantity account (and by extension, other process views of causation), the counterfactual account, and the manipulation account all agree, to a great extent, on the existence of a fundamental level causal relation that I have called causal influence. Furthermore, they pretty much agree on what causal influence relations there are, that is, which concrete events causally influence which other concrete events, though of course they disagree on the metaphysical basis of the relations. If I am correct in holding that any of these variants of the causal influence relation provides enough fine-grained causal asymmetry between the right things to found a causal solution to the problem of explanatory asymmetry, then there is no need to endorse a particular metaphysics of causation in order to commit to a causal account of explanatory asymmetry (though you may have to reject certain other metaphysical theses about causation, as explained below).

I would like to carry over this causal ecumenism to the task of constructing a scientific account of explanation. That is, I would like to advocate a causal account of explanation that requires, from the metaphysics of causation, something that a broad range of theories of causation can provide, and on which they are extensionally if not intensionally agreed. That something will be the causal influence relation.

The causal influence relation in the raw provides a rather unsophisticated account of causal explanation, for reasons examined in depth in chapter two. Thus investigators of causal explanation have tended to appeal in their work to more elaborate metaphysical machinery, such as high level counterfactual dependence or manipulability relations. In so doing, they lose the ecumenical advantage that I hope to retain, that is, the advantage of giving a causal account of explanation that is not held hostage to the details of some particular metaphysics of causation. They are in any case wrong to try to solve all of their explanatory problems metaphysically, for reasons adumbrated in section 1.1—but ecumenism will be my ostensible motivation for now.

Let me put the advantages of the ecumenical approach on display, by pointing to some gnarly metaphysical questions about causation that are left happily unresolved by the minimal assumptions I propose to make about the causal influence relation. First is the question whether causal relations exist only at the fundamental level or whether they exist at all levels, as multilevel accounts assert. My ecumenism entails ignoring, for explanatory purposes, the wealth of high level causal relations offered by multilevel accounts of causation, but it need not mean denying the existence of these relations or their importance to endeavors other than explanation. (They clearly have heuristic use in explanation itself; see section 3.8.) It is quite consistent, then, to endorse (a) a causal approach to explanation, (b) a multilevel account of causation, and (c) the view that only the fundamental level causal relations play a part in explanation—though I suppose that few would choose to do so (see section 3.23).

Second is the question of whether the causal relation is reducible, and if so, to what. On Woodward's view, causal relations—and so, a fortiori, relations of causal influence—are irreducible. On Lewis's view, causal relations are ultimately reducible to a "Humean supervenience basis": facts about causal relations are nothing more than patterns of particular matters of fact (albeit global patterns). On Dowe's view, causal relations are reducible to non-causal matters of fact about persistence and the exchange of conserved quantities, but these physical matters of fact are not necessarily further reducible to Humean facts. My use of the causal influence relation is compatible with any of these positions.

Third is the question of the role of laws in causation. For some writers, following Davidson, two events are causally related only if they fall under a law: *c* is a cause of *e* only if there is a law *All Fs are Gs*, where *c* is an *F* and *e* is an *G*. Others have imposed the even stronger requirement that there must be a law of the form *All Cs are Es*, that is, a law under which the events fall in virtue of the very properties by which they are specified in the causal claim, though this is now thought to be obviously unreasonable. Finally, it is sometimes said that there need be no law at all. What I assume concerning the causal influence relation is consistent with any of these views. (You will observe, however, that the three accounts of causation described above all envisage a close relation between the fundamental physical laws and the causal influence relation—Lewis's and Woodward's because of the central role played by laws in the evaluation of counterfactuals.)

Two further questions left open by my ecumenical assumptions concern the relation between causation, time, and locality. Can there be backwards causation? Must all causation be local? I assume no particular answers to these questions, though the particular accounts of causation mentioned above in some cases do.

All religions, however ecumenical, need someone to burn. And there is, not surprisingly, a view of the nature of causation that is ruled out by even my ecumenism. I have in mind an emergent view of causation, on which there are no causal relations at the level of fundamental physics, only at higher levels. To abandon this possibility is no loss for the proponent of the causal approach to explanation, however, since it is the one view of causation that clearly cannot solve the problem of explanatory asymmetry. The reason: there are explanatory asymmetries in many fundamental physical explanations, such as the case of Rutherford's gold foil (section 1.42). To explain these asymmetries causally, you need a fundamental level causal relation. Causal emergentism has no place in the causal approach to explanation.

This last point serves to reiterate an underlying advantage of ecumenism: what I assume in the way of causal relations is not only something that causal metaphysicians can agree on; it is also just what is needed, and no more, to provide a causal resolution of the explanatory asymmetries.

1.47 Alternatives to the Causal Account

It seems only fair to consider two non-causal approaches to the problem of explanatory asymmetry. The first non-causal approach, due to Hempel (1965a, 374–5) and Salmon (1970), is to find some asymmetry in cases like that of the barometer and the storm that is not causal, but which can be used, like the causal asymmetry, to impose a preferred direction on explanation. Both Hempel and Salmon appeal to the fact that the front that (intuitively) causes the storm will be more reliably correlated with the storm than will the barometer reading (Hempel using a different example). Salmon's more sophisticated approach uses Reichenbach's notion of *screening off*. Roughly, one potential explainer *c* screens off another potential explainer *d* from the explanandum *e* just in case the correlation between *d* and *e* is entirely due, in a non-causal sense spelled out in the formal definition, to a correlation between *c* and *e*, but not vice-versa (Reichenbach 1956).

Reichenbach introduced the notion of screening off, however, to impart his empiricist account of the *causal relation* with the requisite asymmetry. Inevitably, I think, proposals like Hempel's and Salmon's have come to be regarded, not as attempts to provide an asymmetric relation to stand as an alternative to the causal relation, but rather as attempts to provide an empiricist version of the causal relation, in the manner of Reichenbach. When the proposal is regarded in this way, it becomes an attempt to provide an empiricist foundation for a causal account of explanation (as in Salmon 1984), rather than as a genuine alternative to the causal approach to explanatory asymmetry.

Taking a second non-causal approach to explanatory asymmetry, Kitcher argues that the unification account of explanation is well able to handle cases such as that of the flagpole and the shadow (Kitcher 1981). His treatment appeals to features of the unification approach present in any winner-takes-all pattern subsumption view.

On such a view, a phenomenon is not explained by just any pattern under which it can be subsumed. Rather, one particular pattern is uniquely explanatory, namely, the subsuming pattern which maximizes generality, accuracy, and cohesion. (On Kitcher's account, simplicity is also a desideratum, and optimization is global.) Call this pattern the *optimal subsuming pattern* for an explanandum. Now, according to Kitcher, the optimal subsuming pattern for the shadow's length is the pattern described by the law relating flagpole height and shadow length. Thus the flagpole's height explains the shadow's length.

This pattern is not, however, the optimal subsuming pattern for the flagpole's height. There is a pattern that attains a superior mix of accuracy, cohesion, and—in particular—of generality, namely, the pattern corresponding to what Kitcher calls an *origin and development* derivation, a kind of narrative in which the formation and subsequent modification of an object are recounted. All such stories, Kitcher suggests, have something in common, and thus instantiate a single pattern of great generality. It is this origin and development pattern that explains the height of the flagpole.

At this point, Kitcher attempts a grand metaphysical inversion: our notion of causation is, he proposes, derivative of our notion of explanation. One thing is causally relevant to another just in case it is explanatorily relevant, with explanatory relevance being determined by the unification account. Thus Kitcher endorses the thesis that explanatory asymmetry and causal asymmetry go together, but in an unconventional way: rather than relieving the burden on Kitcher's account of explanation by taking care of some potential counterexamples, the association of the two symmetries increases the burden by requiring that all features of the causal relation be derived from the acausal unification criteria for explanation.

By taking on such a heavy load, Kitcher has vastly increased the potential rewards of the unification account's success. But does it succeed? I am not convinced that the origin and development derivations form a real pattern. What do the stories of the building of a flagpole and the formation of the sun have in common? Nothing, except that they are processes of causal production, a feature to which Kitcher cannot appeal. To put the point more formally, any model that could generate both a flagpole and a class G star would, I think, be rather incohesive—so incohesive as to easily nullify its great generality.

1.5 The Pursuit of Explanation

The philosophical study of scientific explanation can be divided into two parts, a description of actual scientific explanatory practice and an evaluation of this practice. The evaluative project in turn has three parts. First, the internal consistency of the practice can be assessed. Second, the assumptions, both metaphysical and physical, implicit in the practice can be assessed. Under this heading you might ask, for example, whether the asymmetry in the causal influence relation is consistent with modern physics. Third, the value of the practice can be assessed. Under this heading you might ask whether there is anything worthwhile, intrinsically or otherwise, in the practice of enquiring into the causes of phenomena (Strevens 2007). This study will be almost entirely concerned with the descriptive project; the second and third evaluative questions will be raised only in passing. (On the third, see sections 4.4 and 7.36.)

It is extremely valuable, I think, to treat the descriptive part of the study of explanation separately from the evaluative part. We should first ask: what do we regard as explanatory? And only then ask: how might we be mistaken? Mixing the descriptive and the prescriptive elements of the study can result in a certain abuse: prescriptive arguments are used to rationalize deficiencies in the description. That is, where an account of explanation fails to capture fully our actual explanatory practice, a normative argument is introduced to the effect that the part of the practice not captured is worthless. Of course,