

Accommodating Explanatory Pluralism

Christopher Pincock (chrispincock@gmail.com)

November 10, 2016 – 9771 words

Abstract: Explanatory pluralism is the view that explanations come in two or more different types. This paper clarifies two versions of explanatory pluralism and considers two very different attempts to make sense of it. On the one hand, an ontic approach isolates genuine explanations only by appeal to facts that obtain in the world. The most promising way for an ontic approach to accommodate explanatory pluralism is to posit different sorts of objective dependence relations. On the other hand, an epistemic approach requires that features of agents appear in any analysis of what a genuine explanation is. I argue that there is a version of an ontic account and a version of an epistemic account that do make sense of explanatory pluralism in their own terms. At the same time, pluralism raises problems for each approach that require further consideration.

I. Strong and weak explanatory pluralism

A pluralist about X maintains that X s come in a variety of different types X_1, \dots, X_n .¹ A pluralist about value, for example, points to different types of values and argues for their philosophical significance. My focus in this paper is explanatory pluralism. To be an explanatory pluralist is to insist that explanations come in several types. Many discussions of explanatory pluralism consider only what could be called minimal pluralism. A minimal pluralist insists that a genuine explanation

¹ I am grateful to the editors and several anonymous referees for their helpful comments on an earlier draft of this paper.

with certain virtues cannot be replaced by an explanation of another type with those very same virtues. This is how Lipton argues for a view that he calls “pluralism”. Scientific explanations of various sorts are needed, and so there is no point, e.g., in trying to reduce all causal explanations to micro-causal explanations: “A good scientific explanation sometimes requires macro causes, sometimes micro causes, and sometimes a combination of the two. When it comes to scientific explanation, we should be pluralists” (2008: 124). Lipton’s argument for pluralism supposes that these different types of explanations tend to realize different explanatory virtues such as strict necessity as opposed to “generality and unification” (2008: 122). What makes a scientific explanation of one sort good is often just not something that can be matched by an explanation of another sort. I take this minimal form of pluralism to be very plausible, and also easy to accommodate on a wide range of views about explanation.

A more controversial form of pluralism claims that for each genuine explanation E_1 of one type there simply is no genuine explanation E_2 of another type that incorporates, subsumes or absorbs E_1 .² For one explanation to absorb another is for that explanation to have the other one as a part. Schematically, if E_1 takes the form of C standing in relation R to E, then it will be absorbed by E_2 when E_2 takes the form of C standing in relation R to E along with other facts, such as that D stands in

² Cf. Reutlinger (forthcoming). He argues that the pluralist must show that there is no theory that covers all explanations. I believe that this places an unfair burden on the pluralist as they must argue that explanations of different types resist any unified theoretical treatment.

relation R to C. Exactly what this comes to depends on whether one adopts an ontic or an epistemic approach to explanation. An ontic approach identifies both the object of the explanation and the explanation itself with facts. What makes some facts explain another fact is a feature of the world as it is independent of human agents. By contrast, an epistemic approach adds an essential reference to human agents and their knowledge states. So in order to say what makes some facts explain another fact, an epistemic view will add additional tests tied to the states of the agents doing the explaining.

Explanatory pluralism requires that explanations come in different types. On an ontic interpretation, what this means is that there is an explanation E_1 of type T_1 of object of explanation O , and the facts making up E_1 are not a part of any more encompassing explanation of any other type.³ An epistemic approach will say something quite similar except this approach can use knowledge states as well to block one explanation from being absorbed into another. As I will discuss in the next section, one type of explanation is causal explanation. So the explanatory pluralist is committed to there being explanations that are not part of any causal explanation. But each type of explanation may have interesting internal relations. For example, one causal explanation may be subsumed under another causal explanation.

On both an ontic and epistemic view, a genuine explanation will require facts that bear the right relation to the fact being explained, and each of these facts will typically be represented by a true proposition. Two sorts of non-minimal explanatory pluralism are examined in this paper. Strong explanatory pluralism

³ I suppose here that O is some fact.

maintains that some explanatory targets have genuine explanations of different types. That is, for some object of explanation O, both E explains O and F explains O and these explanations are of different types. There are two ways to show that alleged explanations of different types are actually of the same type.⁴ Either argue that one explanation actually includes the other or that both are included in a third more encompassing explanation. Consider, for example, two causal explanations of an event. If some light bulb turned on because an electrical current was running through a circuit, then that constitutes one causal explanation for why a light bulb went on. But another explanation of the same type is that a switch was flipped, and allowed the current to run through the circuit, and this turned the light on. This second explanation subsumes the first explanation, and this shows that they are of the same type. There are also cases of genuine explanations of the same target where neither includes the other, but both are subsumed by some third explanation. That one switch was flipped explains why at least one light bulb went on and that another switch was flipped also explains why at least one light bulb went on. But that the department head ordered that more lights be turned on explains both why the first switch was flipped and why the second switch was flipped, and so why at least one light bulb went on. This shows that these two explanations of that target are of the same type.⁵

⁴ These are two sufficient conditions for being of the same type. Necessary and sufficient conditions for being of the same type are given in section II.

⁵ Brigandt (2013) deploys a similar contrast between strong and weak explanatory pluralism. His argument for strong explanatory pluralism concerns explanatory

For strong explanatory pluralism to be true there must be distinct types of explanation. In the next section I introduce three types of scientific explanation: causal, constitutive and abstract. A causal explanation cites the causes of the phenomenon being explained, while a constitutive explanation indicates what composes the phenomenon and how that composition makes the phenomenon obtain. In addition, I argue that there is a third type of explanation that I call abstract. An abstract explanation points to certain abstract characteristics of the system that make the system have certain features. If these are all genuine explanations, and they apply to the very same target, then strong explanatory pluralism is vindicated. There will be explanations of some target phenomenon that are free-standing of one another in the sense that there is no potential to absorb any two of them into some more encompassing explanation. An explanation of a given type, when it is found, provides something that no explanation of any other type can offer.

Strong explanatory pluralism can be contrasted with a weaker explanatory pluralism that merely insists that explanations come in different types. Weak explanatory pluralism does not require that there is some single target that is explained by explanations of different types. It is consistent with this possibility, but also consistent with each type of explanation having its own special sort of

models that “make jointly incompatible idealizations (necessitated by different explanatory aims)” tied to different research programs (2013: 88). This is not the argument I develop here, but I must reserve engaging with this argument for future work. See also Woody (2015) and Potochnik (2015) for related arguments.

explanatory target. For example, one might think that there is a special sort of explanation found in pure mathematics. The object of these explanations is the truth of some mathematical theorem. A purely mathematical explanation of the truth of some theorem might involve a proof that has special characteristics that distinguish it from other proofs that merely show that the theorem is true. One could believe in this type of explanation and yet remain a weak explanatory pluralist. This position would insist that there are no purely mathematical explanations of non-mathematical targets. There is thus no overlap between the objects of these mathematical explanations and the other types of explanation, such as causal explanations. A strong explanatory pluralist denies that the objects of explanations are sorted into these disjoint families. Again, there are some targets of genuine explanations that have two or more types of explanation.

Both the weak and the strong explanatory pluralist face a general challenge that arises for any form of pluralism. Suppose we have a list of different types of explanations such as causal, constitutive and abstract. The pluralist then faces an unappealing dilemma. Either the members of this list have nothing in common or they have something in common. If the members of this list have nothing in common, then it is hard to say why they are actually types of explanation. They may be something more generic such as facts, but they lack any common core that unites them all as explanations. However, if the members of the list do have something in common, and if this is to illuminate how they are all types of explanation, then it is not clear what kind of pluralism can be maintained. A weak pluralist points to mathematical explanations of mathematical theorems and causal explanations of

physical events, and supposes that they are all explanations despite their different targets. The strong pluralist adds that some causal explanations are explanations of the very same things as some constitutive or abstract explanations. Either way, it remains unclear how all these accounts can be explanations and yet fall into irreducibly different types. The pluralist owes us a discussion of what all explanations have in common and what nevertheless divides these explanations with this common feature into distinct types. Otherwise the common feature threatens to unify explanations into a single type and pluralism of any form is blocked.

In the rest of this paper I argue for three claims. First, the diversity of explanations found in scientific practice mandates some form of explanatory pluralism. Second, the most promising form of explanatory pluralism is a version of weak explanatory pluralism that insists that the target of each explanation is a contrast of the form P rather than Q. Third, this flavor of explanatory pluralism fits with a version of an ontic approach and a version of an epistemic approach, but both views face challenges. The ontic approach has difficulty making sense of contrastive facts. The epistemic view can make sense of the explanation of contrasts by appeal to the knowledge states of agents. But it remains unclear how either approach can vindicate the value that scientists place on finding explanations as opposed to merely true descriptions of phenomena.

II. Three types of explanation

Cases drawn from scientific practice can be used to motivate explanatory pluralism. Here I will sketch three cases that support the conclusion that

explanations come in at least three types. These are causal explanation, constitutive explanation and abstract explanation. A case that illustrates these three types is the board of directors of some organization that is made up of people, all of whom are bald.⁶ Suppose that we aim to explain why all the directors are bald. There are three kinds of explanation of this general fact. As we will see, constitutive explanations and abstract explanations are different from causal explanations in virtue of containing a special sort of non-causal relation. A constitutive explanation makes essential use of part/whole relations. An abstract explanation uses the fact that one thing instantiates another. For example, an abstract geometrical structure may be instantiated by a physical system, but the structure is not a part of the system. By contrast, a causal explanation exploits only causal relations. Using these assumptions, I will argue that causal, constitutive and abstract explanations are of different types. The distinctive non-causal relations found in constitutive and abstract explanations block any attempt to subsume them under causal explanations. For similar reasons, we can neither subsume a constitutive explanation under an abstract explanation nor subsume an abstract explanation under a constitutive explanation. A necessary and sufficient condition for being of the same explanatory type, then, is that two explanations exploit the same explanatory relations. If explanation A uses relation R and explanation B uses relation S, then A and B are of different types.

⁶ This case is emphasized for different purposes in Hempel 1965. As will become clear, my treatment of this example is influenced by the classic discussion of Garfinkel 1981 and the more recent Haslanger 2016.

This way of dividing up explanations into types is further motivated by the widely accepted point that adding more facts can spoil an explanation. Suppose, for example, that A stands in relation R to B, and that this fact is a causal explanation of B. It does not follow that the combined fact that C stands in relation S to B and that A stands in relation R to B is also an explanation of B. This “non-monotonic” aspect of explanation holds even when it is the case that the fact that C stands in relation S to B alone is an explanation of B. Combining explanations need not preserve there being an explanation.

One genuine explanation of the fact that the board of directors are all bald is the votes of the membership that elected each director. In a series of elections, first A got the most votes, then B got the most votes, and so on until all the elections are covered. If we add that A is bald, B is bald, and so on until each director is mentioned, we have an explanation of why all of these directors are bald. On many views of causal explanation, this amounts to a genuine causal explanation. Here I suppose that Woodward has developed an adequate account of causal explanation, and our sketch certainly counts as a causal explanation by Woodward’s lights (Woodward 2003). Woodward emphasizes the need to say how the actual situation would have differed if at least one parameter is varied, while others are held fixed at their actual values. Woodward adds the restriction that a parameter is varied by an “intervention”. This limits his test to cases where a causal relation obtains. In the board of directors case, a change in the votes during the election that actually elected A would have resulted in the election of a rival candidate Z. If we suppose that Z is not bald, then this change in the votes would have made it the case that

some of the board members are not bald. For Woodward, this amounts to a causal explanation of why all the board members are bald.

A second explanation notes that A is bald because he lacks sufficiently many hairs on his head. This second explanation would point to a similar condition for B and the other directors. The distinctive feature of this explanation is that it cites the composition of A, B, and the rest in the sense that the lack of hairs are parts of these people. If A was composed differently, and as a result had hairs on his head, then he would not be bald. And if he were not bald, then it would not be the case that all the members of the board were bald. When an explanation appeals to the parts of the phenomenon being explained, then I will call it a constitutive explanation.

My argument that this explanation is of a different type than any causal explanation is that this explanation deploys the part/whole relation in an ineliminable way. So, no causal explanation can fully absorb this constitutive explanation. However, it is not immediately clear that this argument works. It might seem that Woodward's notion of an intervention is flexible enough to accommodate whatever is genuinely explanatory in this explanation. If so, then the explanatory role of the part/whole relation is minimal.

Recent discussions of Woodward's notion of an intervention have highlighted this issue in connection with cases where there are non-causal dependencies between variables (Shapiro & Sober 2007, Woodward 2015). In Woodward's example, a person's level of cholesterol TC is the sum of their LD and HD levels of cholesterol. Woodward claims that TC and LD stand in a non-causal relation of "definitional dependence" (2015: 327). He uses this relation to understand other

non-causal relations of dependence, esp. supervenience relations. In the cholesterol case there is no “relevant” intervention on LD that fixes TC at its actual value. In our case, we can suppose that a person’s baldness B is a variable with values 1 for “bald” and 0 for “not bald”, and also that the density D of the hairs on their head determines their baldness. If D is greater than some threshold, then $B = 1$. If D is below that threshold, then $B = 0$.⁷ But the value of D constitutes the person’s baldness, rather than causing it. An explanation that proceeds through this sort of link is thus quite different than an ordinary causal explanation. The part/whole relation is not eliminated or replaced by wholly causal relations. In this sense, then, my original argument stands.

A third type of explanation of the baldness of the board of directors is available. This is the structural, or what I will call “abstract”, explanation. Suppose that the elections occur in a highly sexist society that gives men many more opportunities for professional advancement. This sexism structures the election of the board members in such a way that it nearly guarantees that all the board members are men of a certain age. If we suppose also that baldness is much more common among men of that age than among women or younger men, then we have a distinct structural explanation of the makeup of the board of directors. There are abstract features of the whole organization and the society that it is a part of that are highly conducive to this outcome.⁸

⁷ Here I ignore the complications associated with the vagueness of this predicate.

⁸ This is not the same as Jackson and Petit’s notion of program explanation. See Pincock 2015: 871-874 for a discussion of the differences.

The special feature of this abstract explanation is that it abstracts away from the constitutive features of the board members. There is nothing special about A, according to this explanation, that made him get elected to the board. For if A had been sidelined through some personal misfortune, and not had the opportunity to run for the board, then the abstract structure of the whole system is such that another candidate A' would have run in his place. And given the character of this system A' is overwhelmingly likely to have been an older man. This shows the gap between our constitutive explanation and our structural explanation. No appeals are made to the particular elements of the system or their internal constitution.

It remains debatable, of course, how different the structural explanation is from any causal explanation.⁹ One might, for example, propose a Woodward-style interpretation of the structure of the society in question. It is certainly true that if this sexist structure were changed to a more egalitarian structure, then the composition of the board with respect to its baldness is highly likely to be changed. What is less clear is whether this sort of change should be thought of as a change in the value of a variable. In the constitutive case we saw that it was important to recognize two sorts of links between variables, namely causal and constitutive. In

⁹ Although Haslanger draws attention to the importance of structural explanations and interprets them in terms of the instantiation of abstract structures, she also appears to view them as a special kind of causal explanation. In particular, Haslanger relates her structural explanations to Dretske's "structuring causes" (2016: 120).

the structural case, the structural features serve as a kind of background against which causal links between variables are established. The structure is thus taken for granted in a way that distinguishes it from any particular causal variable. This suggests that we should contrast what the structure enables or inhibits with what happens within that structure. A change in structure can have dramatic effects that will complicate any notion of the intervention on a structure.

The point can be further supported by appeal to the instantiation relation. An abstract structure is instantiated by a more concrete system. In an abstract explanation the instantiation relation plays a part in the explanation. In our baldness case, there is a complex social structure that maps out the abstract network of gender relations in our society. This structure is instantiated in our society. This fact forms a central part of the abstract explanation for why all the board members are bald. The instantiation relation here cannot be replaced by causal or constitutive relations as this abstract structure neither causes nor constitutes the network of gender relations found in our society. I conclude that this abstract explanation cannot be absorbed into an explanation of either of the other two types.

Our discussion of the board of directors case shows the need for at least the types of explanations that I have called causal, constitutive and abstract. Two more mathematical cases can be used to make the same point (Pincock 2007, 2015). The residents of Königsberg wondered why they had failed to make a circuit of their city that involved crossing each of its seven bridges exactly once. A causal explanation of this pattern could appeal to each of the attempted circuits and indicate how it had

failed by either crossing a bridge more than once or missing a bridge. Each failure has its cause, and this cause can be given a Woodward-style analysis. For example, Wilhelm's failure to complete a circuit arises when he crosses the western most bridge twice. He would not have failed this way if he had turned left rather than right at one point in his journey.¹⁰ A constitutive explanation could appeal to the material that made up the bridges and the people making the crossings. Wilhelm would not have failed in the way he did if he had been paralyzed. Finally, an abstract explanation could appeal to the structure of the bridges. This structure ensures that no attempted circuit would be successful.

An even more mathematical example concerns the laws for how soap-film surfaces meet in stable soap-film configurations. Plateau noticed certain patterns to these meetings that he codified into three laws. A causal explanation of this pattern would indicate the mechanism through which these systems minimize their surface area, subject to the constraints imposed. A constitutive explanation could summarize the spatial arrangement of the parts of each such system and show how they conform to Plateau's laws. Finally, an abstract explanation would show how the patterns found by Plateau follow from a more general mathematical structure. Any instance of that mathematical structure would conform to Plateau's laws.¹¹

¹⁰ One might worry that this causal explanation does not explain the very same fact as the constitutive and abstract explanations. I develop this point in section IV using contrastive facts.

¹¹ See Pincock (2015), Saatsi (2016) and Baron et al. (forthcoming) for more discussion of mathematical explanations of physical phenomena. Andersen

III. Ontic accounts

Causal, constitutive and abstract explanations are different types of explanations.¹² It looks like the same fact is being explained across types and so our cases appear to support what I have called strong explanatory pluralism. Ontic accounts that identify explanations with facts have great difficulty in accommodating strong explanatory pluralism. In the remainder of this section I will consider two ontic attempts to accommodate this kind of pluralism. The first attempt generalizes Woodward's notion of an intervention to cover all three types. The second attempt deploys the concept of ontological dependence to make sense of each of these explanations. Both attempts face the same problem. They wind up with such a weak common feature among explanations that they lose a substantial account of what makes explanations valuable. For this reason, these proposals cannot distinguish explanations from non-explanations.

We have already seen that Woodward's notion of a causal relation tied to interventions is too narrow to include constitutive part/whole relations. The same point holds for structural instantiation relations, as Woodward notes in passing (2003: 220). However, one could try to identify a more generic notion of "difference

 (forthcoming) develops a very different picture of these cases. She uses a notion of a model "holding of" a system to motivate strong explanatory pluralism. I unfortunately lack the space to discuss this important argument here.

¹² These types of explanation have some affinity to Aristotle's efficient, material and formal causes, respectively. I defer to future work an investigation of a modern analogue of Aristotelian final causes in the explanation of human action.

making” that includes all three of these explanatory relations. Woodward himself talks of “what if things had been different”. It might seem that a broader modal test could identify what our three explanatory relations had in common. But this common feature would not undermine explanatory pluralism as the more specific characteristics of these relations could still play a role in individuating types of explanations.¹³

For the board of directors case, the causal explanation meets Woodward’s more demanding intervention test: there is an intervention on the variable that reflects the vote that elected A such that Z is elected instead. This change results in a change in the baldness state of the board, as we supposed that Z is not bald. By passing this more demanding test, the causal explanation also passes a more generic modal test: it tells us how things would have been different, namely how the baldness state would have changed if the vote had gone that way. So far, so good. A similar pattern obtains for the constitutive explanation. Now we explain the baldness of the board via the composition of its members and their internal constitution. The part/whole relation here does not pass Woodward’s intervention test, but it does pass the more generic modal test: if A had been constituted differently, so that A was not an older male, but was instead a woman, then A would not have been bald. So the baldness state of the board would have changed if A’s internal constitution had been changed. Finally, consider the structural explanation that appeals to the instantiation of a sexist social structure. If the system had not instantiated this structure, but instead instantiated the structure of an egalitarian

¹³ See esp. Saatsi & Pexton (2013), Rice (2015), and Reutlinger (forthcoming).

society, then the board would no longer have its baldness state. The structural explanation also indicates what would have been different, but now via its instantiation relation.

The current proposal, then, is that each type of explanation explains by deploying a relation that indicates how things would have been different if various changes had been introduced into the actual board of directors system. What varies across types is *how* this relation gives this modal information. That is why there is a genuine form of pluralism. But there is still a unified core to this class of genuine explanations: if modal information is provided, then one has a genuine explanation.

One problem with this proposal is that it is too flexible.¹⁴ There are simply too many cases where an account that fails to be a genuine explanation deploys a relation that provides the right kind of modal information. Many of these cases can be found in classic objections to Hempel's D-N account of explanation. Consider, for example, the attempt to explain E using C where there is no causal link from C to E, and yet C and E are highly correlated due to some common cause F. Thunderstorms are caused, in part, by a drop in atmospheric pressure. And a drop in atmospheric pressure also causes a barometer to show a lower reading. This generates a strong correlation between a barometer showing a lower reading and a thunderstorm occurring. If a scientist proposed that the barometer's lower reading explained the thunderstorm, then this proposed explanation would be rejected as not genuine.

¹⁴ Another worry is that it fails for cases that involve pure mathematics. See Baron et al. (forthcoming) for a recent discussion. I am grateful to an anonymous referee for emphasizing this problem.

However, this proposed explanation certainly does convey the right kind of modal information. It says how things would have been different: if the barometer had not given the lower reading, then the thunderstorm would not have occurred. This shows that merely conveying modal information is not sufficient for providing a genuine explanation.

Another proposal along these lines is to require that the modal information be conveyed by appeal to one of the following relations: (i) causal, (ii) constitutive or (iii) structural instantiation.¹⁵ The proposed barometer explanation fails this more demanding test because that account did not link the barometer to the thunderstorm by a causal, constitutive or structural relation. This revised modal proposal faces two problems. First, it does not clarify why it is *these* three relations that are needed for an explanation. If a new relation was considered as a supplement to this list, then how are we to tell that it could or could not generate genuine explanations? If providing modal information is not sufficient, it is unclear why providing modal information by one or the other of these relations is sufficient. Second, there are counterexamples like the barometer case that provide modal information via one of these relations, but yet are not genuine explanations.

Consider, for example, a failed constitutive explanation of the board of directors' baldness. It may be the case that any alteration of a board members' genetic makeup that is sufficient to lower their risk of heart attack would also lower their baldness.

¹⁵ A modal approach could of course be developed in other ways. I look forward in particular to Reutlinger's forthcoming work as he clearly recognizes the worry raised in the last paragraph (forthcoming: 8).

So, we can truly say that were some board member to have a lower risk of heart attack, then the board would not be composed entirely of bald people. This proposed explanation conveys modal information by appeal to a constitutive relation that obtains in the actual board, and yet it is not a genuine explanation. If this strategy is to accommodate explanatory pluralism, then a tighter set of conditions must be imposed.

A modal strategy tries to accommodate explanatory pluralism by tying each genuine explanation to a modal fact. A distinct ontic strategy is to focus instead on relations of ontological dependence. As emphasized by Fine, Koslicki and others, ontological dependence relations may obtain even in the absence of the usual modal facts. The set whose only member is the number 3, for example, may be said to ontologically depend on the number 3 despite the necessary existence of both the set and the number 3. So it might seem promising to ground a form of explanatory pluralism on the obtaining of an ontological dependence relation. This is Koslicki's suggestion in her paper "Varieties of Ontological Dependence":

[...] an explanation, when successful, captures or represents [...] an underlying real-world relation of dependence of some sort which obtains among the phenomenon cited in the explanation in question [...] If this connection between explanation and dependence generalizes, then we would expect relations of ontological dependence to give rise to explanations within the realm of ontology, in the sense that a successful ontological explanation captures or gives expression to an underlying real-world relation of ontological dependence of some sort (Koslicki 2012: 212-213).

There is thus a list of dependence relations that includes (i) causal, (ii) constitutive and (iii) structural instantiation. A genuine explanation of E in terms of C involves linking C to E by one of these dependence relations. This dependence need not involve any modal information, and so the presence or absence of modal features is not decisive in the evaluation of the proposed explanation. Instead, what is decisive is whether or not this special sort of relation obtains. Our causal explanation explains by citing the causal relation between the vote and A's presence on the board. The constitutive explanation explains via the constitutive relation that obtains between A's hairs and A's baldness. Finally, the structural explanation functions by appeal to the instantiation relation that obtains between the abstract sexist structure and the society which instantiates it.

One worry about the dependence proposal is that it is hard to figure out what all these ontological dependence relations have in common. One suggestion is

(*) that E ontologically depends on C just is that C makes E obtain.

This natural suggestion faces an overdetermination problem if we add the suppositions that there are distinct types of dependence relation and only one way for something to be made to obtain. Consider, again, the fact that all the members of the board of directors are bald. On the dependence proposal, this fact is explained in three different ways tied up with causal, constitutive and structural dependence. Using (*), if the baldness state depends on its causes, then these causes together make the baldness state obtain. But equally, via (*), if the baldness state depends on its composition, then its composition makes the baldness state obtain. A similar point holds for the structural instantiation relation. The problem now is that there

are three different types of facts, each of which serves to make the baldness fact obtain. How can this be? The dependence proposal must be revised to allow that each dependence relation makes a fact obtain in its own way. There is no competition between these ways and so no risk of overdetermination.

At this point the dependence proposal takes on a somewhat mysterious aura. Explanations explain because they involve these relations and these relations are significant because they make facts obtain, but each type of relation works differently and so can make a fact obtain in a different way. Again we face the problem of saying why certain relations make the list of dependence relations while others are excluded. It may just be a metaphysically primitive feature of the world. But if it is just a primitive feature of the world, then this strategy for accommodating explanatory pluralism leaves us with little recourse for resolving debates about explanation. Someone may propose, for example, that in addition to the way that wholes constitutively depend on their parts, there is also a way that parts holistically depend on the wholes they are a part of. This means that there are “holistic” explanations over and above the causal, decompositional and structural explanations already considered. How can an advocate of our revised dependence proposal combat this suggestion or any other suggestion? Partly for this reason, we lose any link to the value that scientists place on having explanations. If we do not understand what makes a relation a dependence relation, then we also lack an understanding of what makes something an explanation. But scientists do value explanations, and so we must hope that there is some feature that all explanations

have in common that makes the quest for explanation coherent. So far we have not found any way to do this consistent with strong explanatory pluralism.¹⁶

IV. Contrastive facts and a viable ontic strategy

An ontic approach to explanation identifies explanations with facts. These facts explain some other fact according to the sort of relation emphasized by that specific flavor of ontic approach. So far we have considered two unsuccessful ontic attempts to making sense of strong explanatory pluralism. A modal approach zeros in on how some facts are responsible for the modal features of some actual fact. A dependence theory instead posits basic dependence relations that connect the facts doing the explaining and the fact being explained. The overly flexible character of the modal approach showed the need for a tighter connection between explanations and their targets. But the dependence approach faced an overdetermination problem and resolved it by supposing various ways that some facts could make a fact obtain.

A shift to weak explanatory pluralism makes room for a resolution of some of these problems. The overdetermination problem arises only because E explains O and F explains the very same O. If we could somehow distinguish the fact that E explains from the fact that F explains, then there would no longer be any obstacle to maintaining that E makes O obtain and also that F makes O' obtain. One promising way to do this for the cases we have considered is to suppose that the object of these

¹⁶ An ontic view of explanation could add on a further account of the cognitive state known as understanding. This appears to be Strevens' strategy for making sense of explanatory pluralism.

explanations is a contrastive fact. In our discussion so far we have operated with the basic contrast between all the members of the board being bald and it not being the case that all the members of the board are bald. But we could add that there is actually a richer space of contrastive facts in play here, and that an explanation of a given type is suited to explain only one kind of contrastive fact. In this way, the objects of these explanations would themselves be sorted into disjoint types that reflect the types of the explanations. On this position, only weak explanatory pluralism obtains as there is no object that is explained by explanations of different types.¹⁷

An ontic approach can sidestep the overdetermination problem, then, by recognizing only weak explanatory pluralism, and one way to do that is to finely individuate the objects of explanation as contrastive facts. The same maneuver can be used to try to address the other worry about ontic approaches to pluralism. This is that there is no account of why certain relations give rise to explanations and others do not. With contrastive facts at its disposal, the ontic approach can use the character of the contrasts to motivate the explanatory relations and distinguish them from the non-explanatory relations.

To see how this might work, let us reconsider the board of directors case, but now in terms of various contrastive facts. We supposed that A got more votes than

¹⁷ Hitchcock (2012) argues for different types of explanation and that the object of each explanation is a contrast. However, he does not claim that each contrast is apt to be explained by at most one type of explanation. He seems to endorse the dependence proposal discussed in the last section. See esp. 2012: 26.

another candidate Z, and that A was bald and Z was not bald. The contrast then is between the board of directors (including A) all being bald rather than the board of directors (including Z) nearly all being bald. This contrast is well explained by the causal explanation that cites the votes in the election that gave A more votes than Z. But, crucially, this very contrast is not explained by either the constitutive or the abstract explanation. The constitutive explanation considered the parts of each of the actual board of directors and indicated how the actual parts gave rise to the baldness of each. This has no bearing on how Z could have become a board member. Similarly, the abstract explanation cited the instantiation of a sexist social structure. This sexist social structure has no tie to the contrast between A and Z being on the board as that structure is being held in place across this contrast.

What related contrasts, then, are apt to be explained by a constitutive or an abstract explanation? Consider the contrast between the board of directors all being bald rather than some of those very board members not being bald. To explain this we cannot cite the votes that elected the actual board members. We must instead consider the internal constitution of some of those board members. Clearly, if A's internal constitution had been different, such that he had more hairs on his head, then he would not be bald. So we see that a constitutive explanation is well-suited to explain this contrast. The contrast, in effect, holds fixed the chain of events leading up to these people being on the board, but requires us to consider changes in the people's internal constitution. This is why a constitutive explanation is appropriate and no causal explanation can succeed.

The abstract explanation is designed to explain the following contrast: the board of directors all being bald rather than being reflective of the rate of baldness of the general population. Let us suppose that 25% of the population is bald. This contrast can be explained by giving some basis for the gap between the 100% baldness of the board and the 25% baldness of the population that the board is drawn from. The fact that the society instantiates a sexist social structure does explain this contrast as it classifies the actual society in a way that shows how the two percentages could diverge so sharply. There is a kind of top-down structuring to the events leading up to these board members all being bald. By contrast, in other societies where a different, more egalitarian social structure is instantiated, more of a match between the population and the board is to be found. Neither the causal explanation nor the constitutive explanation fits this contrast. The causal explanation considers how causes operate within the given social structure and so does not factor in what is due to that structure itself. The constitutive explanation varies only the internal constitution of the actual board members, and so also does not consider the role of the abstract social structure.

Schematically, then, we have three kinds of contrastive facts and we can suppose that there is something about the kind of contrastive fact that makes it well-suited to be explained only by an explanation of a single type. Roughly, when a contrast is tied to a difference that could have been made through causes changing events, while fixing the constitutive character and the broader abstract structure, then a causal explanation is mandated. When a contrast relates to a change in the internal constitution of one or more elements, while not varying the causes between

events or the broader abstract structure, then a constitutive explanation is required. Finally, when a contrast invokes a difference between types of systems, then only an abstract explanation will cite the right kind of factor that is responsible for those differences across systems. Looking to the operations of causes or the internal constitution of the elements of the actual system will fail to make sense of that sort of contrast.¹⁸

An ontic account that embraces this kind of weak explanatory pluralism thus avoids the overdetermination problem and is able to motivate their list of explanatory dependence relations. The relations that figure in explanations naturally fall out of the character of the contrasts being explained. Does this show that there is an ontic route to accommodating explanatory pluralism?

V. A viable epistemic strategy

Lipton has done the most to distinguish contrastive facts from other kinds of facts, and to make contrastive facts the proper objects of many scientific explanations (Lipton 2004, 2008). Lipton notes that we sometimes explain contrastive facts of the form P rather than Q and uses this practice to argue that a contrastive fact is distinct from any other kind of fact. Consider, for example, the attempt to reduce each contrastive fact P rather than Q to a conjunction of P and

¹⁸ Sober 1986 and Hitchcock 2012 independently suggest that contrasts have presuppositions. The character of these presuppositions may explain why only one type of explanation works for a given contrast. However, Sober and Hitchcock focus on causal explanation and do not seem to have extended this insight to non-causal cases.

not-Q. Lipton rejects this reduction by pointing out that some explanations of the contrastive fact fail to be explanations of the conjunction. Here he uses the assumption that an explanation of a conjunction must explain both of the conjuncts. But in a classic case like Jones rather than Smith having paresis, one genuine explanation is that Jones rather than Smith has untreated syphilis. This explanation of the contrast fails to explain the conjunction that Jones has paresis and Smith does not because untreated syphilis rarely leads to paresis. So the contrastive fact is different from the conjunctive fact by a principle of indiscernibility of identicals. One cannot reply to Lipton's argument by saying that this fact is after all a genuine explanation of the conjunction, but just a poor explanation. For even if it is a poor explanation of the conjunction, the fact is a good explanation of the contrast. This shows that the contrastive fact and the conjunctive fact are different.

Lipton also offered a helpful difference condition that gives a necessary condition on a causal explanation of a contrastive fact:

To explain why P rather than Q, we must cite a causal difference between P and not-Q, consisting of a cause of P and the absence of a corresponding event in the case of not-Q (Lipton 2004: 42).

We have seen this principle at work in our causal explanation of the baldness of the board of directors. One cause of the board of directors being bald (with A a member) rather than not bald (with Z a member) is the vote that elected A rather than Z. That vote caused A to be elected, and it corresponds to the absence of Z's getting more votes.

Lipton unfortunately does not generalize his difference condition to non-causal explanations of contrasts. This is somewhat surprising as he argues for the existence of non-causal explanations. We can fill this gap by recasting Lipton's condition in more generic terms. Now we must speak of "explanatorily relevant factors" instead of causes, and also not suppose that these factors pertain just to events. As we have seen, the explanatorily relevant factors may relate to an entity's internal constitution or to the abstract structure instantiated by the system. With this wider range of cases in mind, Lipton's principle becomes:

To explain why P rather than Q, we must cite an explanatorily relevant difference between P and not-Q, consisting of a feature of P and the absence of a corresponding feature in the case of not-Q.

We considered a constitutive explanation of the contrast between all the board of directors being bald rather than those very same members not all being bald. This explanation cited A's internal constitution and its tie to being bald. Similarly, we reviewed the abstract explanation of the contrast between all the board of directors being bald rather than reflecting the rate of baldness of the general population. This explanation turned on the sexist social structure that was instantiated in the actual society. Its presence directed older men, and so people who were highly likely to be bald, to the board, while the presence of a more egalitarian social structure leads to more representative outcomes for boards in other sorts of societies.

It remains unclear how to approach explanations of contrastive facts on an ontic view of explanation. Recall that an ontic view identifies explanations with facts. So even though there are true propositions for each explanation, these

propositions are not integral to the explanation when properly conceived. The difficulty is that the contrastive facts and the explanations we have arrived at are often tied up with the interests of the scientists investigating the world or other contextual factors over and above the facts themselves. The contextual aspects of contrastive explanation are emphasized by van Fraassen and Garfinkel, and even Lipton seems willing to concede their importance:

What makes one piece of information about the causal history of an event explanatory and another not? The short answer is that the causes that explain depend on our interests. But this does not yield a very informative model of explanation unless we can go some way towards spelling out how explanatory interests determine explanatory causes. One natural way to show how interests help us to select among causes is to reveal additional structure in the phenomenon to be explained, structure that varies with interest and that points to particular causes. The idea here is that we can account for the specificity of explanatory answer by revealing the specificity in the explanatory question, where a difference in interest is an interest in explaining different things (Lipton 2004: 33).

Lipton's discussion suggests a world of facts with an overwhelming array of explanatorily relevant factors for each of these facts. An agent investigating the world then selects out a contrast as her object of explanation. Once this selection is made, and only once this selection is made, is there a determinate answer for whether any proposed explanation of this contrast is a genuine explanation of that contrast. This is partly because of Lipton's difference condition. But we could add

that the contrast that has been selected is only able to be explained by one type of explanation. So the selection of the contrast not only cuts down the number of explanatorily relevant factors, but also specifies that only one type of factor is relevant.

The viability of an ontic account of explanation turns on its making sense of this selective function of agents.¹⁹ A non-ontic, more epistemic alternative could match many of the advantages of an ontic account by focusing on questions of knowledge. On the ontic view, the contrastive fact that is the object of explanation is a genuine fact that emerges somewhat mysteriously out of the non-contrastive facts that obtain in a given situation. Whenever P and not-Q obtain, then P rather than Q obtains, although these are different facts. Accommodating weak explanatory pluralism has led the ontic account to privilege these contrastive facts as the objects of many scientific explanations. An epistemic alternative takes a different view of these contrastive facts. On this alternative approach, it is agents who know the conjunctive fact that P and not-Q and this knowledge is then presupposed in any legitimate explanatory question. When an agent knows the conjunctive fact, then they are able to pose the explanatory question “Why P rather than Q?”. However, on

¹⁹ My narrow concern is quite different from Wright’s sweeping attack on ontic approaches. Essentially, Wright assumes that “explaining designates a processual activity, which static or inert objects like sundials are incapable of performing” (2015: 29). But a defender of an ontic approach can and should distinguish the act of explaining from the explanation itself. Similarly, one should distinguish the act of pointing from the object that is pointed out.

the epistemic approach there is no need to posit any further contrastive fact. Instead, it is the agent's knowledge and their interests together that generate a legitimate question. The legitimacy of the question is established by factors beyond the obtaining of the conjunctive fact.

The epistemic alternative is non-ontic because it invokes factors beyond the facts in the world by themselves when determining whether or not something is a genuine explanation. These factors pertain to knowledge states and other states of the agents investigating the world. It is partly in virtue of these factors that an explanatory question is legitimate. One worry that is often raised against this sort of proposal is that it makes the existence of genuine explanations too closely tied to features of agents. As a result, it looks like we must index the genuineness of an explanation to a time, person or research community. Newton had a genuine explanation of the fall of bodies on Earth *for Newton*, while Einstein had a genuine explanation of the fall of bodies on Earth *for Einstein*.

Given what we have seen so far, the epistemic account sketched here is not vulnerable to this form of relativism. For, just as with the ontic account, we can suppose that a contrast is apt to be explained by only one type of explanation. And, with Lipton, we can suppose that which explanations of this type are genuine is fixed only by the contrast and the facts that obtain in the world. Contextual factors like states of knowledge and interests do serve to determine which explanatory questions are legitimate for which agents. But the role of the context is limited to just this step. Once the explanatory question is in place, only certain explanations

count as genuine, and what makes them genuine is that they reflect actual presences and absences of the right sort of explanatory factors.

This epistemic approach can endorse Woodward's picture of the very limited role of "pragmatics" in a theory of scientific explanation: "what we want to explain – the particular explanandum we want to account for – often depends on our interests or on contextual or background factors" (Woodward 2003: 229). Unlike Woodward, though, this epistemic account makes explicit how some knowledge states figure into the selection of a legitimate object of explanation.

We have arrived, then, at two somewhat equally matched strategies for accommodating explanatory pluralism. Both the ontic view and the epistemic view first retreat to weak explanatory pluralism by finely individuating the objects of explanation in terms of contrasts of the form P rather than Q. The ontic view supposes that there is a contrastive fact in the world, and that its internal character makes it apt to be explained by only certain kinds of other facts in the world. The epistemic approach instead adds an account of legitimate explanatory questions. A legitimate question takes the form of "why P rather than Q?" and presupposes the knowledge of P and not-Q. But as with the ontic view, the epistemic view adds that this question selects for certain kinds of explanatorily relevant factors in the world. A genuine explanation will then be an account that picks out some facts that do bear the right kind of relation to the contrastive question. The ontic view claims that all the features of a genuine explanation relate only to facts in the world, and that the characteristics of agents are irrelevant. The epistemic view maintains that the world plays an important role, but that a full account of what makes an explanation

genuine must start with legitimate explanatory questions. Which questions are legitimate will vary with a person's states, especially their states of knowledge and interests. However, on this epistemic view, that is the only role for context and pragmatics.

Each strategy faces its challenges. The ontic view must clarify the nature of contrastive facts and their relationship to non-contrastive facts. The epistemic view needs to flesh out what makes a question legitimate. If all questions are legitimate, then we risk trivializing explanation (Kitcher & Salmon 1987). Either way, the arguments of this paper show that it is not easy to make sense of explanatory pluralism. Whatever strategy turns out to be the best, there are many extant approaches to explanation that fail to accommodate explanatory pluralism while doing justice to the value that scientists place on discovering genuine explanations.

References

- Andersen, H. (forthcoming), 'Complements, Not Competitors: Causal and Mathematical Explanations', *British Journal for the Philosophy of Science*.
- Baron, S., M. Colyvan & D. Ripley (forthcoming), 'How Mathematics Can Make a Difference', *Philosophers' Imprint*.
- Brigandt, I. (2013), 'Explanation in Biology: Reduction, Pluralism, and Explanatory Aims', *Science and Education* 22: 69-91.
- Garfinkel, A. (1981), *Forms of Explanation: Rethinking Questions in Social Theory* (New Haven: Yale University Press).
- Haslanger, Sally (2016), 'What Is A (Social) Structural Explanation?', *Philosophical Studies* 173: 113-130.
- Hempel, C. (1965), *Aspects of Scientific Explanation* (New York: Free Press).
- Hitchcock, C. (2012), 'Contrastive Explanation', in M. Blaauw (ed.), *Contrastivism in Philosophy* (New York: Routledge), 11-34.
- Kitcher, P. & W. Salmon (1987), 'Van Fraassen on Explanation', *Journal of Philosophy* 84: 315-330.
- Koslicki, K. (2012), 'Varieties of Ontological Dependence', in F. Correia and B. Schneider (eds.), *Metaphysical Grounding: Understanding the Structure of Reality* (Cambridge: Cambridge University Press), 186-213.
- Lipton, P. (2004), *Inference to the Best Explanation*, Second edition (New York: Routledge).

- Lipton, P. (2008), 'CP Laws, Reduction and Explanatory Pluralism', in J. Hohwy & J. Kallerstrup, *Being Reduced: New Essays on Reduction, Explanation and Causation* (Oxford: Oxford University Press), 115-125.
- Pincock, C. (2007), 'A Role for Mathematics in the Physical Sciences', *Nous* 41: 253-275.
- Pincock, C. (2015), 'Abstract Explanations in Science', *British Journal for the Philosophy of Science* 66: 857-882.
- Potochnik, A. (2015), 'The Diverse Aims of Science', *Studies in History and Philosophy of Science* 53: 71-80.
- Reutlinger, A. (forthcoming), 'Is There A Monist Theory of Causal and Non-Causal Explanations? The Counterfactual Theory of Scientific Explanation', *Philosophy of Science*.
- Rice, C. (2015), 'Moving Beyond Causes: Optimality Models and Scientific Explanation', *Nous* 49: 589-615.
- Saatsi, J. & M. Pexton (2013), 'Reassessing Woodward's Account of Explanation: Regularities, Counterfactuals, and Noncausal Explanations', *Philosophy of Science* 80: 613-624.
- Saatsi, J. (2016), 'On the 'Indispensable explanatory role' of mathematics', *Mind* 125: 1045-1070.
- Shapiro, L. & E. Sober (2007), 'Epiphenomenalism – The Do's and Don'ts', in G. Wolters & P. Machamer (eds.), *Thinking About Causes: From Greek Philosophy to Modern Physics* (Pittsburgh: University of Pittsburgh Press), 235-264.

Sober, E. (1986), 'Explanatory Presupposition', *Australasian Journal of Philosophy* 64: 143-149.

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

Woodward, J. (2015), 'Interventionism and Causal Exclusion', *Philosophy and Phenomenological Research* 91: 303-347.

Woody, A. (2015), 'Re-orienting Discussions of Scientific Explanation: A Functional Perspective', *Studies in History and Philosophy of Science* 52: 79-87.

Wright, C. (2015), 'The Ontic Conception of Scientific Explanation', *Studies in the History and Philosophy of Science* 54: 20-30.