

Conjunctive Explanations and Inference to the Best Explanation*

Jonah N. Schupbach[†]

Abstract

Fitting with the theme of the special issue, this paper explores one way in which discussions of the nature of scientific explanation can inform the proper statement of explanatory inference. Inference to the Best Explanation (IBE) advises reasoners to infer exactly one explanation. This uniqueness claim apparently binds us when it comes to “conjunctive explanations,” distinct explanations that are nonetheless explanatorily better together than apart. To confront this worry, explanationists qualify their statement of IBE, stipulating that this inference form only adjudicates between competing hypotheses. However, a closer look into the nature of competition reveals problems for this qualified account. Given the most common explication of competition, this qualification artificially and radically constrains IBE’s domain of applicability. Using a more subtle, recent explication of competition, this qualification no longer provides a compelling treatment of conjunctive explanations. In light of these results, I suggest a different strategy for accommodating conjunctive explanations. Instead of modifying the form of IBE, I suggest a new way of thinking about the structure of IBE’s lot of considered hypotheses.

1 The Challenge of Conjunctive Explanation

Sometimes two explanations are better than one. This may happen, for example, in cases of “explanatory pluralism” when theories each do qualitatively different explanatory work. An object’s existence can be explained either by referring to its causes or its function—cf. (Wright, 1976). One hypothesis may explain an event by telling us a causal-mechanical story leading up to the event, while another may perhaps explain the same event by referring to a nomic regularity that

*I owe special thanks to Vincenzo Crupi, David Glass, and attendees at the 11th MuST Conference in Philosophy of Science on “Models of Explanation” (University of Turin, Italy; June 11-13, 2018) for helpful conversations and feedback regarding this project. I am also very grateful for two sources of financial support for this project. While working on this paper, I was supported by the *Charles H. Monson Mid-Career Award* administered by the University of Utah’s Philosophy Department. The research for this publication was also made possible through the support of a grant (#61115) from the John Templeton Foundation. The opinions expressed in this publication are those of the author(s) and do not necessarily reflect the views of the John Templeton Foundation.

[†]University of Utah, Department of Philosophy; 417 CTIHB, 215 S. Central Campus Dr., Salt Lake City UT 84112; +1(801)585-5810 (office); jonah.n.schupbach@utah.edu.

the event instantiates—cf. Salmon’s (1981; 2001) “friendly physicist” example. In such cases, accepting a plurality of explanations provides us with a richer understanding of the explanandum. More generally, several explanations are better than one just when the explanatory benefits of accepting them all outweigh the costs (in complexity and otherwise). In such cases, I will say that the distinct potential explanations in question are “conjunctive”, and I will refer to the above observation as the phenomenon of “conjunctive explanation.”

The observation that there exist conjunctive explanations might seem mundane. But conjunctive explanations apparently spell trouble for Inference to the Best Explanation (IBE), at least in its simplest formulations. In an exchange with Peter Lipton, Wesley Salmon criticizes IBE precisely for mishandling conjunctive explanations. The problem, as Salmon (2001, p. 67) presents it, is with IBE’s “uniqueness claim”: “The phrase, ‘inference to *the best* explanation,’ involves a uniqueness claim that is difficult to justify.” IBE will bar us from inferring truths, gaining richer understanding, accepting otherwise appealing explanatory hypotheses, etc. in conjunctive explanation cases, since it ostensibly mandates that we only choose the single best explanation.

In response to this challenge of conjunctive explanation, Lipton offers what is now widely regarded as a necessary qualification on IBE: “[IBE] is meant to tell us something about how we choose between *competing* explanations: we are to choose the best of these. But among compatible explanations we need not choose” (Lipton, 2001, p. 104; cf., 2004, pp. 62-63). Call this move “Lipton’s hedge.” The suggestion is that Salmon’s criticism holds no sway against IBE once it’s properly qualified. This is because IBE so hedged does not even attempt to adjudicate between non-competitors, the working presumption being that conjunctive explanations cannot compete with one another.

A proper evaluation of Lipton’s hedge must take into account what it means for potential explanations to *compete* with one another. This question results in a balancing act that Lipton and his followers must manage. To give IBE a formulation that does not force us to choose between conjunctive explanations, Lipton’s hedge restricts IBE’s domain of applicability to cases in which the alternative explanations compete. In aiming to rule out the problematic cases (of conjunctive explanation) and only these cases, this competition qualification is susceptible to two potential errors. On the one hand, the account of competition might be too strong, ruling out more than the problematic cases and overly restricting IBE’s domain of applicability. The problem in this case would be that IBE does guide us in reasoning between explanations not considered to be competitors according to such a strong account. On the other hand, to the extent that the proposed account of competition is too weak (not ruling out all of the problematic cases), the challenge of conjunctive explanation remains. In this case, the problem would be that conjunctive explanations can compete in the salient, weaker sense. And we would not want IBE to force us to choose between such conjunctive explanations in such a case any more than in cases where conjunctive explanations don’t compete. The hope for Lipton and his followers then is that there is a plausible explication of competition that strikes the right balance in order to rule out approximately all and only the problematic cases of conjunctive explanation.

In this paper, I will argue that no candidate account of competition manages

to strike the desired balance. A stronger “all-out” reading of competition allows one to bypass the challenge of conjunctive explanation, but only at the expense of absurdly restricting IBE’s domain of applicability. A weaker, more generally palatable explication of competition fails to meet the challenge of conjunctive explanation. This failure of contemporary accounts of competition to strike the desired balance motivates a reassessment of Lipton’s hedge. I argue that Lipton’s hedge was never needed in the first place by suggesting an alternative way of responding to the challenge of conjunctive explanation. The upshot is a defense of IBE, as traditionally formulated, with Lipton’s hedge completely trimmed.

2 All-Out Competition and the Ubiquity of IBE

Lipton’s hedge requires that the explanations being compared in any instance of IBE compete with one another. But what exactly does it take for explanatory hypotheses to compete? In the above quote, Lipton goes along with a popular trend in philosophy of science and assumes (perhaps only for the sake of simplicity) that hypotheses compete only when they are incompatible. Potential explanations may be incompatible by virtue of being directly inconsistent. Mutually exclusive descriptions of flag pole height and position of the sun constitute incompatible, competing potential explanations of the length of the pole’s shadow. But potential explanations may also be rendered incompatible by the evidence they aim to explain. The hypotheses that John committed the robbery and that Bill committed the robbery are compatible, but they may be rendered incompatible by evidence showing that there could only possibly have been one robber acting in the case. Either way, when explanatory hypotheses compete in the extreme sense of being incompatible, they cannot possibly be true together. Accepting either potential explanation accordingly provides us with a decisive case for rejecting the other. Call this extreme notion of competition “all-out competition”—to be contrasted with a less extreme sense of competition in the next section.

With respect to the challenge of conjunctive explanation, it’s easy to see the appeal of explicating competition as all-out. Lipton’s hedge provides a convincing response to this challenge, at least in part, *because* it invokes the extreme reading of competition. If IBE only adjudicates between competing hypotheses, *and competition amounts to incompatibility*, then IBE should manifestly require us to choose at most one hypothesis. Conjunctive explanations pose no challenge at all to IBE, because there is not a situation to which IBE applies in which we ever somehow miss out by only inferring one explanation. The underlying assumption here is that incompatible explanations cannot constitute conjunctive explanations, that it is never explanatorily better to accept an unsatisfiable conjunction of hypotheses. And that surely seems right.

So Lipton’s hedge, when teamed with the all-out notion of competition, rules out all of the problematic cases of conjunctive explanation. But the natural follow-up question to ask is whether it rules out *only* those cases. Are there cases in which IBE helpfully guides us to infer between potential explanations that are not all-out competitors? If so, then Lipton’s hedge overly restricts IBE’s domain of applicability. This question is all the more pressing once one recalls a familiar point commonly made by Lipton and explanationists more generally: that IBE is

ubiquitous, being an extremely useful inference form with an impressively expansive domain of applicability (Lipton, 2004, pp. 1-2). Does IBE lose its intuitively expansive reach in light of Lipton's hedge?

In fact, the answer is a troublingly emphatic and obvious yes. Many (indeed, plausibly *most*) canonical instances of IBE compare potential explanations that are compatible with one another. Indeed, this is true of nearly all of Lipton's own examples of IBE at work.

Lipton's foremost example is the Semmelweis case. Working in the maternity division of the General Hospital in Vienna in the 1840s, Ignaz Semmelweis struggled to explain why three times more women in the first maternity ward were dying of "childbed fever" than in the second ward of the same hospital. Ward one was staffed by medical students, whereas ward two was overseen entirely by midwives. The potential explanations considered and tested by Semmelweis included the following:

- H_1 . The midwives in ward two encouraged women to give birth on their sides, whereas the medical students had women give birth on their backs. The latter birthing position somehow promotes childbed fever.
- H_2 . A priest was more often seen in ward one on his way to administering last rites to dying patients. This has a pernicious psychological influence on birthing women, which subsequently promotes childbed fever.
- H_3 . Unlike the midwives, medical students in ward two were routinely conducting autopsies. Childbed fever is promoted by an infection of "cadaveric matter" from the hands of such students.

Eventually, Semmelweis famously inferred H_3 as the best explanation of his accumulating evidence.

Lipton takes this to be a paradigmatic example of IBE at work, regularly drawing upon this example to develop his account of the nature and power of explanatory inference. What Lipton does not seem to recognize is that his use of this example clashes with his response to the challenge of conjunctive explanation. For if IBE is only meant to "tell us something about how we choose between *competing* explanations," and competition is cashed out as "incompatibility" then Lipton's favorite example of IBE at work is in fact not an example of IBE at all. After all, H_1 , H_2 , and H_3 are manifestly compatible. Any combination of these hypotheses could have been true prior or posterior to considerations of Semmelweis's collected evidence. If Lipton is right that IBE guides reasoning in the Semmelweis case, then he is wrong that IBE only adjudicates between incompatible explanations.

Upon further reflection, it's plausible that such cases—in which IBE adjudicates between compatible alternatives—are more the norm than the exception. Explanationists commonly draw instances of IBE from such contexts as detective work, historical science, medical diagnosis, and diagnostic settings more generally (e.g., diagnosing the failure of a car engine from the observable "symptoms"). Unless explanatory hypotheses from such contexts are intentionally framed so as to exclude one another, it is straightforward to think of such cases as often comparing compatible potential explanations.

Let us emphasize the point by glossing over a couple more examples. Lip-

ton (2001, pp. 95-96) writes, “When a detective infers that it was Moriarty who committed the crime, he does so because this hypothesis would best explain the fingerprints, blood stains and other forensic evidence [...] Moriarty’s guilt would provide a better explanation of the evidence than would anyone else’s.” Lipton uses this example to demonstrate the fallibilistic nature of IBE; some other potential explanation than the best may turn out to be the actual explanation. But this example can just as well be used to demonstrate the fact that it is possible, in very typical cases, for more than one of the potential explanations compared by IBE to be an actual explanation—i.e., that more than one of these turn out true, given their joint satisfiability. In this particular example, the crime may after all have been committed by Moriarty and someone else.

Scientists debate the explanation of the mass extinction at the Cretaceous-Paleogene (K-Pg) boundary—which included the mass extinction of the dinosaurs about 66 million years ago. Common potential explanations include bolide impact, massive outbreaks of volcanic activity, climate change, continental drift and sea level regression, and so on. Notably, while it is perhaps more common to look for a “smoking gun” amongst these alternatives (Cleland, 2011, p. 554), many scientists opt instead for inferring some combination of these alternatives—e.g., (Archibald et al., 2010). The scientists involved in this debate do not think of the various alternatives as incompatible; instead they argue either that one alone suffices as the best explanation of the evidence, or that more than one of the compatible alternatives should indeed be inferred—see (Schupbach and Glass, 2017) for further discussion of this example. Accordingly, this is another case that the explanationist will be keen to describe as potentially involving IBE, despite the fact it involves reasoning between recognizably compatible alternatives.

In sum, if Lipton is right that IBE only provides a model of inference between incompatible explanations, then he is wrong to think of the above (and any number of other such examples) as instances of IBE. But he’s not wrong about that; these are paradigm examples of IBE at work, instances of IBE if anything is. And so Lipton must be wrong in thinking that IBE only properly provides a model of inference between incompatible explanations.

3 Lipton’s Hedge Refined

Lipton’s hedge, when combined with the all-out explication of competition provides a convincing response to the challenge of conjunctive explanation, but only by absurdly restricting IBE’s domain of applicability. To pin the blame immediately on Lipton’s hedge itself however would be too quick. It may be that Lipton is correct to think that IBE provides a model of inference only between competing explanations, and that he only goes astray when he explicates competition as all-out. That is, maybe the problem is not with the hedge per se, but with Lipton’s identification of competition with incompatibility.

Recent work by Schupbach and Glass (S&G) proves helpful here. S&G (2017) argue that competition is not plausibly explicated as mutual exclusivity. As they claim and demonstrate through examples, it is simply too easy to think of actual cases from various contexts of human reasoning in which recognizably compatible hypotheses are thought of and inferentially treated as competitors. S&G thus

offer a probabilistic explication of what it takes for hypotheses to compete with one another in light of a body of evidence (as well as a formal measure of the degree to which hypotheses compete with one another apropos some body of evidence). Most importantly for our present purposes, their explication allows for cases in which compatible hypotheses nonetheless compete.

S&G motivate their account of competition by suggesting that a “mutual exclusivity” account falls short of a satisfactory, general account of competition in no less than two ways. The first is that it implies that competition is all-or-nothing (this objection applies equally well to the more general incompatibility explication that Lipton invokes). But hypotheses may compete by disconfirming each other to varying degrees without fully precluding one another. Consider the following simple variation on Lipton’s detective case:

Moriarty and Smith, v1. Moriarty and Smith are both house burglars working in the same area, but they are also sworn enemies who are extremely unlikely to ever burgle together. Bob reports to the police that his front window has been broken and that all of his valuable belongings are missing from the house.

A detective investigating the case may rightly view H_S : *Smith burgled the house* and H_M : *Moriarty burgled the house* as distinct (but compatible) explanations of the reported evidence. Moreover, given the background information about Moriarty and Smith’s relationship, the detective might rightly view these potential explanations as competing *to the extent that* H_S and H_M disconfirm each other. To make sense of such cases, S&G require that hypothesis competition be accounted for gradationally, as a matter of degree.

Noting that hypotheses may compete to varying degrees helps shed light on the nature of competition in cases where hypotheses disconfirm one another, but to some less than maximal extent—as in Moriarty and Smith, v1. However, in many actual cases, competing hypotheses may not even disconfirm one another directly (i.e., prior to consideration of the explanandum). This is plausibly the case in at least two of the examples we have described above. In the Semmelweis example, it is not at all clear that birthing position’s having an influence on childbed fever rates would somehow lower the probability of a priest’s presence also having such an influence. This is unclear, but it’s clear that these hypotheses compete in this case. The nature of their competition with one another must come down to something other than a disconfirmatory (probability-lowering) relation between them then. Similarly, in the K-Pg extinction case, far from disconfirming one another (does volcanic activity decrease the chance of bolide impact?!), some of these historical hypotheses may even *confirm* one another to some extent. Nonetheless, many scientists persist in viewing these as competitors.

These observations point to another shortcoming with the mutual exclusivity account. It neglects an important sense in which hypotheses can compete—indirectly, via the relevant body of evidence E . Consider another variation on Lipton’s detective:

Moriarty and Smith, v2. Moriarty and Smith often rob houses together in well-informed, carefully planned ways. They would never knowingly rob the

police chief's house. Moriarty knows that 123 Main is Bob's house but doesn't know that Bob is the police chief; Smith knows Bob is police chief but doesn't know where he lives. E includes this information as well as Bob's recent discovery that his front window has been broken and that all of his valuable belongings are missing from the house.

While H_S and H_M may confirm one another in general, the detective may rightly consider them as competitors in this particular case; relative to *this* body of evidence E , it's unlikely they collaborated.

In cases of distinctively indirect competition, adopting either hypothesis undermines any support E provides for the other. This may be because the evidence itself places otherwise mutually confirming (or independent) hypotheses in a relation of mutual disconfirmation with one another. Or it may be because the evidence is fully accounted for by one of the hypotheses alone, in which case no support from E accrues any longer for the other hypothesis. In order to make sense of competition in cases like this, S&G require that an appropriate account of competition accommodate two distinct paths to hypothesis competition: a direct path and an indirect path via the evidence.

S&G develop a confirmation-theoretic measure of the "net" degree (taking into account both the direct and indirect paths) to which hypotheses H and H' compete with one another relative to a particular body of evidence E . The full official statement of this measure is somewhat complex, but S&G prove a theorem that simplifies matters. They show that net degree of competition is formally equivalent to average degree of disconfirmation conditional on E . Using the log-likelihood measure of confirmation, the degree to which a proposition ϕ disconfirms another ψ (conditional on a proposition χ) is measured as the degree to which ϕ confirms $\neg\psi$ (conditional on χ):

$$C_l(\phi, \neg\psi|\chi) = \log \frac{P(\phi|\neg\psi \wedge \chi)}{P(\phi|\psi \wedge \chi)}.$$

Thus, the net degree to which hypotheses H and H' compete with one another relative to a particular body of evidence E can be represented as follows:

$$\begin{aligned} \text{Comp}(H', H/E) &= [C_l(H, \neg H'|E) + C_l(H', \neg H|E)]/2 \\ &= \left[\log \frac{P(H|\neg H' \wedge E)}{P(H|H' \wedge E)} + \log \frac{P(H'|\neg H \wedge E)}{P(H'|H \wedge E)} \right] /2. \end{aligned}$$

Moreover, S&G explicate qualitative judgments of hypothesis competition as positive degree of net competition; i.e., the judgment that H and H' compete with one another relative to E is explicated using the inequality $\text{Comp}(H', H/E) > 0$. Because net competition can (as above) be represented as average degree of disconfirmation, we may equally well explicate the judgment that H and H' compete with one another relative to E as an assessed positive degree of disconfirmation, using the inequality $C_l(H, \neg H'|E) > 0$ (or using $C_l(H', \neg H|E) > 0$, since these imply one another). Or we may represent this judgment probabilistically as $P(H \wedge H'|E) < P(H|E) \times P(H'|E)$. Finally—and this will prove the most useful statement of all—Glass (2012, Theorem 1) proves that the following inequality is

yet another equivalent condition for qualitative competition:

$$\log \left[\frac{P(E|H \wedge H')P(E|\neg H \wedge \neg H')}{P(E|H \wedge \neg H')P(E|\neg H \wedge H')} \right] + \log \left[\frac{P(H|H')P(\neg H|\neg H')}{P(H|\neg H')P(\neg H|H')} \right] < 0. \quad (1)$$

If we accept Lipton's hedge, along with S&G's explication of competition, then IBE is no longer so absurdly restricted in scope. The resulting qualified inference form again guides us in Lipton's paradigmatic cases. Recall that in both the Semmelweis case and the K-Pg extinction case (and potentially also in the detective case, depending on the details), we take the potential explanations involved to be at once perfectly compatible apart from the explanandum but to compete relative to this evidence. The sense in which they compete is indirect. But additionally, note that they compete indirectly not because E introduces an incompatibility between them, but simply because any one of these explanations arguably suffices on its own to account for E . In light of E , reason compels us to choose between these potential explanations for the simple reason that accepting more than one would arguably be epistemically overblown.

S&G's account implies that this can be a genuine source of competition; this is easiest to see in terms of the last formal statement of competition. Condition (1) involves the sum of two terms, which can roughly be seen as respectively explicating the notions of direct and indirect degrees of competition.¹ Examining these two summands carefully sheds formal light on some exact paths to competition. Most relevant to our current purposes, note that the second summand—corresponding to the notion of direct competition between H and H' —is effectively a wash when H and H' are approximately independent of each other. Focusing on the first summand in such cases then, H and H' will be deemed competitors with respect to E when the denominator is greater than the numerator, $P(E|H \wedge \neg H') \times P(E|\neg H \wedge H') > P(E|H \wedge H') \times P(E|\neg H \wedge \neg H')$.

Notably, this inequality can easily attain in cases where one, but really *only* one, of the hypotheses is needed to account for the evidence. When this is true, $P(E|H \wedge \neg H')$ and $P(E|\neg H \wedge H')$ will be quite high. In fact, even if both hypotheses account for the evidence somewhat better than either individually, it will still be the case that $P(E|H \wedge \neg H') \approx P(E|H \wedge H')$ and $P(E|\neg H \wedge H') \approx P(E|\neg H \wedge \neg H')$. By contrast, since one or the other hypothesis is needed to account for E in the envisioned scenario, $P(E|\neg H \wedge \neg H') \ll 1$. The upshot is that H and H' compete with respect to E in such a scenario, since $P(E|H \wedge \neg H') \times P(E|\neg H \wedge H') \gg P(E|H \wedge H') \times P(E|\neg H \wedge \neg H')$. This is a prime example of a scenario in which two hypotheses can strongly compete with one another (relative to some E) despite the fact that they are otherwise entirely compatible with—possibly even mutually supportive of—one another!

Qualified in this way, IBE requires us to choose between explanatory hypotheses so long as $Comp(H', H/E) > 0$. The new question is whether the proper balance is now struck. IBE is not absurdly restricted in its domain of applicability, but is it overly restrictive in what it allows us to infer? Unfortunately, it turns out

¹Though only roughly. S&G themselves explicate degrees of direct and indirect competition separately, and they prove that net degree of competition is a simple sum of these two component measures. Unfortunately, there is no neat correspondence between S&G's component measures and the two individual summands in condition (1).

that the challenge of conjunctive explanation comes back for revenge at this point; there are hosts of cases in which it is problematic for IBE to require us to choose between hypotheses that compete on S&G's generalized account.

4 The Challenge of Conjunctive Explanation 2.0

Lipton's hedge, combined with the all-out reading of competition, restricted IBE to cases in which it indeed is always appropriate to infer at most one competing hypothesis. But IBE, so restricted, lost much of its intuitively vast domain of applicability. S&G's construal of competition provides IBE with a wider, more intuitive domain of applicability. But it reopens the door to cases in which it is undesirable for IBE to keep us from inferring more than one of the alternative hypotheses.

To see this, first consider a final variation on Lipton's detective case:

Moriarty and Smith, v3. Moriarty and Smith are far and away the busiest and most notorious house burglars working in a town. They are also sworn enemies who are extremely unlikely to ever burgle together. Moriarty always leaves an "M" pendant at the scenes of his crimes, while Smith's trademark is to leave the water running into a plugged sink, flooding the houses he robs. Bob reports to the police that his valuables are missing from his now-flooded house, where he has also discovered the familiar "M" pendant.

A detective examining this case may rightly be strongly compelled to accept both H_S : *Smith burgled the house* and H_M : *Moriarty burgled the house*. Given the evidence of the case E , this certainly seems to be a better option than accepting H_S or H_M alone. Two explanations are better than one here, making this a case of conjunctive explanation. Importantly however, these explanations are not better together because they mutually support one another but because they are both separately supported by the evidence. E strongly confirms both H_S and competitor H_M individually, but it does nothing to unify them. H_S and H_M disconfirm one another unconditionally *and conditional on E*.

Accordingly, S&G's account validates the judgment that these hypotheses compete. Applying condition (1) to this case, we have the following qualitative criterion for competition between H_S and H_M :

$$\log \left[\frac{P(E|H_S \wedge H_M)P(E|\neg H_S \wedge \neg H_M)}{P(E|H_S \wedge \neg H_M)P(E|\neg H_S \wedge H_M)} \right] + \log \left[\frac{P(H_S|H_M)P(\neg H_S|\neg H_M)}{P(H_S|\neg H_M)P(\neg H_S|H_M)} \right] < 0$$

Regarding the first summand, this case is meant to inspire the following judgments. The evidence is by far best accounted for by the conjunction $H_S \wedge H_M$. However, it is partially accounted for (made somewhat probable) by either hypothesis taken alone (conjoined with the negation of the hypothesis). Since these two criminals are far and away the most active burglars working in the area (and given the presence in this case of their trademarks), the evidence remains very unlikely indeed if neither criminal was at work. Probabilistically, we may explicate these judgments simply using the following inequalities: $P(E|H_S \wedge H_M) > P(E|H_S \wedge \neg H_M)$ [$P(E|\neg H_S \wedge H_M)$] $> P(E|\neg H_S \wedge \neg H_M)$. These inequalities fail




Figure 1: $\ln [P(E|H_S \wedge H_M)P(E|\neg H_S \wedge \neg H_M)/P(E|H_S \wedge \neg H_M)P(E|\neg H_S \wedge H_M)]$ plotted as a function of $y = P(E|\neg H_S \wedge H_M) = P(E|H_S \wedge \neg H_M)$ and $x = P(E|\neg H_S \wedge \neg H_M)$ —with $P(E|H_S \wedge H_M)$ fixed at .95. Values in the displayed range tend to be negative (thus contributing to the fact that H_S and H_M compete), and they are never strongly positive.

to determine whether the first summand above is positive or negative, but they plausibly suggest that the term is negative or at least not strongly positive (Figure 1).

Now consider the second summand. Given that they are sworn enemies, “extremely unlikely to ever burgle together,” each criminal is much more likely to have burgled a house generally if the other suspect did not. The assumption that either criminal did in fact burgle a house thus makes it more likely that the other did not (than did). Alternatively, given the same considerations along with the suggestion that these burglars are collectively (though not typically conjointly) responsible for the vast majority of burglaries in the area, the assumption that either criminal did not burgle the house makes it more probable that the other criminal did (than did not). Probabilistically, $P(\neg H_S|H_M) \gg P(H_S|H_M)$ and $P(H_S|\neg H_M) > P(\neg H_S|\neg H_M)$. The upshot is that the second summand cannot be negative, and the details of the case suggest that it is substantially so. It is thus primarily because of this aspect of the case that we may safely conclude, on S&G’s account, that H_S and H_M compete.

The fact that one would want to accept both H_S and H_M in light of E is evidently not because they do not compete. That is, this is not intended to be a counterexample to S&G’s account of competition. It is right to think of these hypotheses as being (potentially strongly) in competition with one another with respect to this evidence for the simple reason that they (potentially strongly) disconfirm one another—i.e., lower each other’s respective probabilities—in the light of this evidence. Any justification accruing to $H_S \wedge H_M$ in this case is accounted for solely by way of E ’s providing strong support for H_S and H_M individually, and this *despite* the fact that each hypothesis goes some non-negligible way to rebutting the other hypothesis conditional on E (and unconditionally). The example thus serves to show that there are cases in which it would be explanatorily

better to accept multiple, *competing* explanations of E ; competing hypotheses can provide conjunctive explanations.

Guided by the above example, it is not difficult to characterize formally an entire family of such examples in which it can be explanatorily better to accept multiple competing explanations than to choose between them. The following jointly satisfiable², probabilistically explicable conditions are characteristic of such examples:

- C1. H and H' compete with one another with respect to E : $Comp(H', H/E) > 0$, and so $P(H \wedge H'|E) < P(H|E)P(H'|E)$.
- C2. E confirms each hypothesis individually, conditional on the other: $P(H|E \wedge H') > P(H|H')$ and $P(H'|E \wedge H) > P(H'|H)$.
- C3. H and H' together account for the evidence better than either does individually: $P(E|H \wedge H') > P(E|H \wedge \neg H')$ and $P(E|H \wedge H') > P(E|\neg H \wedge H')$.

H and H' may plausibly provide conjunctive explanations in such cases, at least when E is *explanatorily* better accounted for by the conjunction $H \wedge H'$ than by either hypothesis taken alone. That is, allowing that the inequality in likelihoods used to explicate C3 may be achieved apart from explanatory considerations and contexts, still certain explanatory virtues plausibly have their logical effect via just such an inequality in likelihoods—e.g., power (Schupbach and Sprenger, 2011; Schupbach, 2017). But then the problem of conjunctive explanation is once again a serious problem, even for Lipton's hedged version of IBE. H and H' compete with one another *apropos* E , but it can be explanatorily better to accept both than to choose between them. IBE is problematic if it forces us to choose between them. We want to be able to infer the conjunction $H \wedge H'$ as “the” best explanation of E in cases like this, even if they compete.

The skeptical reader might wonder at this point whether it actually might be *epistemically* worse to do what is *explanatorily* better in these cases. That is, one might wonder whether the conjunction $H \wedge H'$ is overall worse off than $H \wedge \neg H'$ or $\neg H \wedge H'$, even if it happens to account for E better in the sense of making E more likely. If the answer is ‘yes’, then perhaps we really do want IBE to force us to choose between H and H' in such cases, even if this may mean not doing what is explanatorily best (though IBE would surely be in need of a new name in this case). Thus, it is very much worth highlighting the subset of cases demonstrating the consistency of C1, C2, and C3 with the following:

- C4. H and H' are overall more plausible together in light of the evidence than either is alone (i.e., conjoined with the negation of the other): $P(H \wedge H'|E) > P(H \wedge \neg H'|E)$ and $P(H \wedge H'|E) > P(\neg H \wedge H'|E)$.

Cases satisfying C1-C4 are ones in which $H \wedge H'$ may have not only explanatory considerations, but overall net epistemic considerations in its favor—at least assuming a Bayesian perspective from which net epistemic value is measured by posterior probabilities. An inference rule is surely problematic if it precludes us

²In all cases, satisfiability was established using Fitelson's (2008) decision procedure PrSAT as implemented in his corresponding *Mathematica* package, available at <http://fitelson.org/PrSAT/>.

from even considering the possibility of inferring $H \wedge H'$ in such cases. But that's just what Lipton's hedge does, when combined with S&G's generalized explication of competition.

5 Trimming Lipton's Hedge

Let's take stock. It can be explanatorily best to accept multiple distinct explanatory hypotheses. Such conjunctive explanations seem to pose a serious challenge to IBE, since this inference form ostensibly guides us to infer *the single* best explanation of our explanandum. Responding to this challenge, Lipton's idea was to qualify IBE by restricting it to cases in which we are comparing competing explanations. However, given Lipton's own interpretation of the notion of competition, this qualification greatly *over*-restricts IBE's domain of applicability, absurdly barring from the ranks many (perhaps most) of IBE's canonical instances. Thankfully, a subtler, generalized explication of competition does not seem to lead to this consequence; however, it results in a hedged version of IBE that again fails to meet the challenge of conjunctive explanation.

Explanations that compete, according to this generalized account, thereby each provide some (possibly strong) reason against accepting the other(s). However, we have highlighted the possibility that this reason may be outweighed by the explanandum's relation to each candidate explanation individually so that the conjunction (of competitors) indeed provides the overall best explanatory account of the explanandum. In other words, when explanations compete, there is some reason not to accept both. But what ultimately matters is whether the explanatory payoff outweighs the cost of accepting competitors. If there are net explanatory gains to accepting multiple, distinct explanations, then IBE should allow us to accept multiple explanations, regardless of whether they compete.

Attempts to meet the challenge of conjunctive explanation by hedging IBE such that it only adjudicates between competing explanations do not alas appear to be successful. Does this mean that the challenge of conjunctive explanation is devastating for IBE? That may be the moral of this story if a competition hedge were the only plausible way to respond to this challenge. But I want to suggest that Lipton's hedge was ultimately a distraction to IBE research. The final consideration offered in the previous paragraph suggests that the question of whether explanations compete is ultimately orthogonal to the central question of what is explanatorily best. This observation inspires a straightforward response to the challenge of conjunctive explanation that makes no mention of the notion of competition.

Indeed, I suggest the simplest (some would say naive) formulations of IBE can be understood as already providing a response to this challenge. The "challenge" of conjunctive explanation is thus only challenging for those who do not understand the statement of IBE in the way I propose. The central question is how to understand the phrase "the best explanation". It is common to point to one source of ambiguity in this phrase: the various, distinct dimensions along which humans tend to evaluate the explanatory goodness of hypotheses (Schupbach, 2017, p. 41). However, another ambiguity underlies the means by which IBE handles conjunctive explanations. On one reading, "the best explanation" might

refer to the single most explanatory hypothesis. On another reading, “the best explanation” might refer to the inference that is explanatorily best, i.e. the most explanatory conclusion. The phenomenon of conjunctive explanation teaches us that these two readings are indeed importantly different: the explanatorily best conclusion might involve inferring more than one explanatory hypothesis.

Conjunctive explanations pose no threat to IBE when “the best explanation” is understood in the second sense. If it is overall explanatorily best to infer multiple distinct explanations (i.e., explanatory hypotheses), then that is exactly what IBE will guide us to do when “the best explanation” refers to the conclusion or inferential move that is overall explanatorily best.

It is tempting to think that the word “explanation” is, or at least ought to be, always associated with a single explanatory hypothesis (however hypotheses may be individuated) rather than with an explanatory inference, move, or conclusion. It is certainly natural to use it in this way; this is in fact how I have largely been using the word throughout this paper. If one insists that this is the only proper usage, then the phrase “the best explanation” no longer is ambiguous in the required way. However, the same straightforward response to the challenge of conjunctive explanation plausibly remains available to IBE’s defenders, with only minor adjustments to the statement of IBE. “Inference to the Best Explanation(s)” or the clunkier “Inference to the Most Explanatory Conclusion” suggest themselves as options.

That said, “explanation” and “the best explanation” do seem to be ambiguous in the required way, allowing for the suggested interpretation. Recall the K-Pg extinction case. As noted, scientists involved in this debate do not think of the various alternative hypotheses as incompatible; but they also do not think of them as necessarily being part of distinct explanations. While some of these scientists hold out hope that one hypothesis alone suffices as the best explanation of the evidence, many contemporary scientists accept more than one of the alternatives as jointly constitutive of *the* fullest explanation. No clash arises here between the inference of multiple explanatory hypotheses and IBE’s uniqueness claim, if the uniquely best explanation (i.e., most explanatory conclusion) can take the form of the conjunction of hypotheses under consideration.

This last statement of the point raises one final important question: if “the best explanation” may refer to a conjunction of individual hypotheses, can it also refer to other Boolean combinations of these? More generally, if we understand IBE in the proposed way, then how should we think about the structure of IBE’s lot of available, potential explanations? Plausibly, there is no principled ban on the sorts of compound hypotheses we may consider for inference. Any form that may provide superior explanatory goodness is up for consideration. If conjunctions are allowed, why not disjunctions? Why not indeed? If a detective only knows that a house has been broken into in a neighborhood where Smith and Moriarty typically work alone, but are collectively responsible for the vast majority of break-ins, the best explanation might in fact be the disjunction $H_S \vee H_M$. To commit to anything more specific than this would be to stretch the explanation beyond what the evidence and background information of the case allow. But if disjunctions are allowed, then why not material conditionals? And can’t denials (in the form of negations) sometimes provide potential explanations?

If “the best explanation” is associated with the most explanatory individual hypothesis, then it’s natural to think of the lot of considered potential explanations as simply being the set of individual hypotheses. By contrast, the present proposal amounts to thinking of the lot of potential explanations as the set containing these considered hypotheses along with their Boolean combinations. What matters is which combination of considered hypotheses best explains the explanandum, not what logical form the various options take. Regarding cases of conjunctive explanation, this move allows reasoners to infer more than one individual hypothesis when that is the explanatorily best option, all the while leaving IBE’s domain of applicability entirely open.

Note that this proposal does not amount to offering a different qualification on IBE in the place of Lipton’s hedge. The point of this response is that there is a way of understanding the simplest, unqualified formulations of IBE so that conjunctive explanations don’t pose a challenge at all. The so-called challenge of conjunctive explanation does not point to a weakness in IBE properly construed; it rather betrays a misunderstanding of the most defensible statement of IBE. As such, there was never a need for Lipton’s hedge in the first place. And that’s a good thing, since this paper goes some way toward showing that this hedge ultimately proves unworkable.

References

- Archibald et al., J. D. (2010). Cretaceous extinctions: Multiple causes. *Science*, 328:973.
- Cleland, C. E. (2011). Prediction and explanation in historical natural science. *British Journal for the Philosophy of Science*, 62:551–582.
- Fitelson, B. (2008). A decision procedure for probability calculus with applications. *The Review of Symbolic Logic*, 1(1):111–125.
- Glass, D. H. (2012). Can evidence for design be explained away? In Chandler, J. and Harrison, V. S., editors, *Probability in the Philosophy of Religion*, pages 79–102. Oxford University Press, Oxford.
- Lipton, P. (2001). Is explanation a guide to inference? a reply to Wesley C. Salmon. In Hon, G. and Rakover, S. S., editors, *Explanation: Theoretical Approaches and Applications*, pages 93–120. Kluwer Academic, Dordrecht.
- Lipton, P. (2004). *Inference to the Best Explanation*. Routledge, New York, NY, 2nd edition.
- Salmon, W. C. (1981). Rational prediction. *British Journal for the Philosophy of Science*, 32(2):115–125.
- Salmon, W. C. (2001). Explanation and confirmation: A Bayesian critique of Inference to the Best Explanation. In Hon, G. and Rakover, S. S., editors, *Explanation: Theoretical Approaches and Applications*, pages 61–91. Kluwer Academic, Dordrecht.

- Schupbach, J. N. (2017). Inference to the Best Explanation, cleaned up and made respectable. In McCain, K. and Poston, T., editors, *Best Explanations: New Essays on Inference to the Best Explanation*, pages 39–61. Oxford University Press, Oxford.
- Schupbach, J. N. and Glass, D. H. (2017). Hypothesis competition beyond mutual exclusivity. *Philosophy of Science*, 84(5):810–824.
- Schupbach, J. N. and Sprenger, J. (2011). The logic of explanatory power. *Philosophy of Science*, 78(1):105–127.
- Wright, L. (1976). *Teleological Explanations*. University of California Press, Berkeley.