

Should Explanations Omit the Details?

Penultimate draft of a paper to appear in the

British Journal for the Philosophy of Science

Abstract

There is a widely shared belief that the higher level sciences can provide better explanations than lower level sciences. But there is little agreement about exactly why this is so. It is often suggested that higher level explanations are better because they omit details. I will argue instead that the preference for higher level explanations is just a special case of our general preference for informative, logically strong, beliefs. I argue that our preference for informative beliefs entirely accounts for why higher level explanations are sometimes better – and sometimes worse – than lower level explanations. The result is a step in the direction of the unity of science hypothesis.

1. Introduction
2. Background: Is Omitting Details an Explanatory Virtue?
 - 2.1. Anti-reductionist arguments
 - 2.2. Reductionist argument
 - 2.3 Logical strength
3. Bases, Links and Logical Strength
4. Functionalism and Fodor's Argument
5. Two Generalizations
6. Should the Base Really be Maximally Strong?
7. Anti-Reductionist Arguments Regarding the Base
8. Should the Antecedent of the Link Really be Maximally Weak?

1. Introduction

It is widely believed that explanations can be improved by omitting details. But why is it good to omit details? One might think that it can be good to omit details because it can be good to provide *less* information. I will argue to the contrary – explanations are always improved by adding information. And sometimes we can *increase* the information provided by omitting details. This looks paradoxical, but happens in the *antecedent of a conditional*; omitting details from the antecedent of a conditional logically strengthens the conditional, providing us with more information.

In the background is the hypothesis of the unity of science (Carnap 1934, Nagel 1951, Oppenheim and Putnam 1958). Let ‘the higher level sciences’ refer to all sciences other than fundamental physics. One part of the hypothesis of the unity of science is that fundamental physics provides better explanations than higher level science. An important line of objection draws on examples to argue that omitting details improves an explanation, and that higher level explanations are better because they omit such details (Putnam 1967, Fodor 1974). I will argue that the wrong moral has been drawn from these examples. They show that *logically strong* explanations are better; they don’t show that higher level explanations are better. Thus, one objection to the unity of science hypothesis is undermined.

Part 2 explains the background reductionist and anti-reductionist arguments; part 3 contains the core argument, introducing the distinction between links and bases and applying it to explanations that omit details; part 4 applies the distinction to functionalism, arguing that functional laws are logically weak, and therefore provide poor explanations; part 5 generalizes the account along two dimensions; part 6 develops the view that the base should be logically strong; part 7 defends the view that the base should be logically strong from Garfinkel’s (1981) objections; part 8 defends the view that the link should be logically strong; part 9 concludes.

2. Background: Is Omitting Details an Explanatory Virtue?

The topic of this paper is explanatory reduction. We’ll focus on the explanation of events, setting aside explanatory reduction of laws. Let’s assume we have a token event that is predicted with the same

probability by both a lower level and a higher level explanation. Can we say anything in general about which explanation is better? *Anti-reductionists* favour higher level explanations; *reductionists* favour lower level explanations.¹ In this section I will explain the opposing arguments, then suggest how they can be reconciled.

2.1. *Anti-reductionist arguments*

The driving force behind many anti-reductionist arguments are well-known examples:

Putnam's Peg (Putnam 1967)

Explanandum: A square peg with 1 inch sides fails to go through a round hole with a diameter of 1 inch.

Higher level explanation: The peg is square

Lower level explanation: A description of the position of every molecule of the peg and hole.

Garfinkel's Rabbit (Garfinkel 1981)

Explanandum: Rabbit r gets eaten.

Higher level explanation: The fox population is high

Lower level explanation: Rabbit r passed through the capture space of fox f.

Yablo's Pigeon (Yablo 1992)

Explanandum: A pigeon trained to peck at red cards pecks at a scarlet card.

Higher level explanation: The card is red

Lower level explanation: The card is scarlet.

¹ Compare Weslake 2010. He discusses reduction as Fundamentalism: 'that there is no dimension of explanatory depth along which the nonfundamental sciences can provide deeper explanations than those provided by fundamental physics'. p.274 Depth seems to be a place-holder for an explanatory virtue.

The Glass (Jackson and Pettit 1992)

Explanandum: The sealed glass, in which water was heated, cracked.

Higher level explanation: Water was boiling

Lower level explanation: Molecule A hit the glass with momentum greater than z.

The Conductor (Jackson and Pettit 1992)

Explanandum: The conductor was annoyed.

Higher level explanation: Someone coughed

Lower level explanation: Bob coughed.

Thermodynamics (Strevens 2014)²

Explanandum: The gas filled the balloon.

Higher level explanation: The gas is at 100 Kelvin

Lower level explanation: The gas consists of a billion molecules moving with velocities V_1 - $V_{1,000,000,000}$

There are differences between these examples that we'll address later; for now, what matters is what they have in common. These examples are often taken to show that explanations should omit details, from which it seems to follow that higher level explanations can be better than lower level explanations³.

² This is the closest I will get to engaging with a specific science. As my thesis is general, it would be distracting to spend too much time engaging with the details of a specific example.

³ See Putnam 1967 p.138, Fodor 1974 p. 103, Garfinkel 1981 p.56, Jackson and Pettit 1992 p.11, and also MacDonald (1992, p. 86, 90–92), Haug (2011, p.1150), Clarke (2016) and Batterman (forthcoming) among others; see Sawyer 2002 for a helpful overview. Technically, it is said we should omit *irrelevant* details. But this qualification doesn't really add anything. Of course we should not omit relevant details. Anyway, I will argue that omitting details is not a fundamental virtue, so relevance plays no role in my account.

But why is it good to omit details? I will briefly survey the literature and argue that no convincing reason has been given to think that omitting details is an *explanatory virtue* (where an explanatory virtue is a good-making feature of an explanation).

First, perhaps details should be omitted when *alternative details* would have produced the same event. Such details would not make a difference to the outcome, and perhaps explanations should only mention features that make a difference to the outcome.⁴

But now we need a specification of what the alternative details are. This suggests a contrastive theory of explanation (in the explanans at least), according to which A rather than B explains E. I will argue in section 7.2 that the thesis that explanations are contrastive does not entail that details should be omitted.

Second, Weslake (2010) suggests the following answer:

Why believe that [omitting details] provides a genuine [virtue]? My central claim...is that [that omitting details is a virtue] provides the best explanation for the truth of [the thesis that higher level explanations can give better explanations].

p.288

But this is circular if the ultimate aim is to explain why higher level explanations are better than lower level explanations. Grant that *if* omitting details is a virtue, then this explains why higher level explanations can give better explanations. But this is no argument that omitting details is a virtue. The reductionist denies that higher level explanations can give better explanations, so has no need of an explanation of it. Furthermore, we have been given no account of why omitting details is a virtue (we've only been told that a preference for higher level explanations can be explained if omitting details is a virtue). And Weslake's account leaves open that we might find a better explanation for the benefits of higher level explanation – indeed this is what I will suggest.

⁴ See Strevens 2004, 2009.

Third, perhaps less detailed explanations can explain more *actual* events. This is related to the virtue of unification.⁵ For example, if we explain why a gas filled a balloon by reference to the exact velocity of the gas's molecules, then this explanation will probably not apply to any other events. By contrast, if we explain why a gas filled a balloon by reference to the gas's temperature, then this explanation will apply to other gas filling events.

But it is also plausible that *more* detailed explanations can explain more actual events. We can even use the same example – the explanation of why a gas filled a balloon by reference to the exact velocity of its molecules explains the precise way in which the balloon expanded i.e. molecule x of the balloon moved to location 1, then molecule y of the balloon moved to location 2 etc. So we cannot simply say that less detailed explanations explain more actual events.

Fourth, perhaps less detailed explanations can explain more *possible* events. Less detailed explanations are less sensitive to the initial conditions, and so are more counterfactually robust i.e. they continue to be explanations across a greater range of possible worlds.⁶ Counterfactual robustness does seem to be an explanatory virtue. In fact, I think it is. But *why* is counterfactual robustness an explanatory virtue?

Perhaps it's a fundamental fact about explanation that counterfactual robustness is a virtue. I have no objection to this view, but we should only posit fundamental facts when necessary. It will emerge that on my account, counterfactual robustness might be a virtue, but is not a fundamental virtue; counterfactual robustness is a good feature of an explanation to the extent that it indicates that the explanation is logically strong. Logical strength is the fundamental virtue.

2.2. Reductionist arguments

⁵ Compare Kitcher [1981]. I'm inclined to say it *is* the virtue of unification, but this might be controversial.

⁶ Compare Skyrms [1980 p.11], Nozick [1981 p.12], Woodward (2003), Woodward and Hitchcock (2003a, 2003b), White (2005), Weslake [2010]

So far I've argued that it's not obvious why omitting details is a virtue of explanations. And there are powerful arguments that adding details improves explanations – details that might only be described using the concepts of a lower level science.⁷ Here are a few examples:

Non-detailed: Socrates died because he drank poison

Detailed: Socrates died because he drank hemlock.

Non-detailed: Bob is crying because she is either happy or sad

Detailed: Bob is crying because she is happy.

Non-detailed: The building collapsed because it was structurally unsound

Detailed: The building collapsed because tension in the concrete keystone caused a crack.

In these cases the detailed explanations seem to be better, suggesting that lower level explanations are better.

Furthermore, the preference for details is emphasized by one strand of thought in the 'new mechanistic' philosophy of science – Machamer, Darden, and Craver 2000, Darden and Craver 2002, Darden 2006 and Craver 2007 emphasize completeness and specificity as an explanatory virtue.⁸ Darden 2007 is perhaps the most explicit:

A mechanism sketch is an incomplete model of a mechanism. It characterizes some parts, activities or features of the mechanism's organization, but it leaves

⁷ For other defences of reductionism see Kim (1992, 1993), Sober (1999) and Dizadji-Bahmani, Frigg and Hartmann (2010).

⁸ See also Churchland 1981, Ylikoski and Kuorikoski 2010. For dissent see Levy & Bechtel (2013).

gaps...[P]roblematically, sometimes [the gaps] are masked by filler terms that give the illusion that the explanation is complete when it is not. p.113

And Elster (1983) writes that 'to explain is to provide a mechanism, to open up the black box and show the nuts and bolts, the cogs and wheels of the internal machinery' (p.23).

So we have a puzzle – sometimes extra details seem to improve an explanation and sometimes fewer details seem to improve an explanation. Michael Strevens (2011) puts his finger on the tension:

On the one hand, we say that an explanation is deep when it goes far down toward the physical level, the level of detail at which ultimate causal underpinnings are found. On the other hand, we also say that an explanation is deep when it has a certain striking generality - when it attributes the phenomenon to be explained not to some very particular set of initial conditions, but to some high-level, abstract, often virtually mathematical state of affairs.

He aptly calls this the Goldilocks problem (Strevens 2008) - that of characterizing what it is for an explanation to be just right in its level of detail. My approach dissolves this problem. Rather than conflicting desiderata, I will argue that there is only one desideratum, and the apparent tension comes from failing to distinguish the two parts of an explanation.

2.3 Logical strength

I suggest we focus on logical strength. My thesis is:

*Logical strength is a fundamental explanatory virtue.*⁹

⁹ Compare Woodward and Hitchcock 2003b p.195

(I'll take logical strength to be equivalent to *informativeness* i.e. 'A is more informative than B' means that A entails B (and B does not entail A).)

Usually adding details increases logical strength; *but in the antecedent of a conditional the usual effects are reversed*. More details means less information. This accounts for the tension. Positing a preference for logical strength explains our judgments about omitting details and higher level explanations.

The issue is complicated by the fact that there are two ways of moving from a lower level to a higher level explanation. One way is to omit details; the other is to use a functional concept.

Example:

Start with the lower level explanation that someone died because they ingested cyanide.

(A) We can generate a higher level explanation by *omitting details* of the chemical, and saying only that ingesting a substance in the cyano group (of which cyanide is a member) caused death.

(B) Alternatively, we could say that ingesting a poison caused death, where 'poison' is understood to refer to anything that causes disturbances in organisms i.e. 'poison' is a *functional* term.

Higher level laws¹⁰ generated by *omitting details* in the antecedent (A) are logically *stronger* than lower level laws; the claim that all substances in the cyano group cause death is *stronger* than the claim that cyanide causes death. Higher level laws generated by using functional concepts (B) are logically *weaker* than lower level laws; the claim that poisons cause death is *weaker* than the claim

¹⁰ I don't hold that explanations need to contain laws, so I should say 'generalization', or 'causal statement', or what I really want to say: 'link'. But I won't introduce this terminology until the next section.

that cyanide causes death. So I will argue that omitting details in the antecedent of the law improves the explanation (section 3) while moving to a functional law worsens the explanation (section 4).

3. Bases, Links and Logical Strength

Explanations have two parts – a basis and a link.¹¹ This division is explicit in Hempel and Oppenheim (1948) for whom

‘the explanans falls into two subclasses;
one of these contains...sentences...which state specific antecedent conditions;
the other is a set of sentences...which represent general laws’ p.137.

The antecedent conditions are the *basis*, and the law is the *link*. It is widely agreed that explanations do not need to appeal to laws; but I will assume we need some kind of conditional to *link* the basis to the thing-to-be-explained.¹²

Now return to my claim that logical strength is a fundamental explanatory virtue. How are conditionals logically strengthened? Conditionals are logically strengthened by *weakening* the antecedent.¹³ For example,

‘if you play well then you will win’

is logically strengthened by weakening the antecedent with a disjunct:

‘if you play well or play above average then you will win’.

So the most informative conditionals have the weakest antecedents i.e. have the fewest details in the antecedents. This is where details should be omitted – in the antecedent of the conditional. And this improves the explanation by making it *more* informative. This is why higher level explanations are sometimes better than lower level explanations.

¹¹ See Schaffer (forthcoming)

¹² Or at least, my account will only apply to explanations containing conditionals.

¹³ And by strengthening the consequent, though this won’t matter here.

Consider Jackson and Pettit's example of the conductor who is annoyed because someone coughed. It is Bob who coughed, but anyone coughing would have annoyed the conductor. Should an explanation of the conductor's annoyance include the fact that it was Bob who coughed? The anti-reductionist says 'omit Bob'; the reductionist says 'include Bob'. How to reconcile these conflicting intuitions? I suggest that we can accommodate both intuitions by noting that the most informative, logically strong, *link* doesn't mention Bob.

Logically weak link: If Bob coughs then the conductor is annoyed

Logically strong link: If someone coughs then the conductor is annoyed

By contrast, when it comes to the *base*, adding the lower level details increases logical strength.

Logically weak basis: Someone coughed

Logically strong basis: Bob coughed

Putting these together, the most informative explanation is:

Logically strong basis: Bob coughed

Logically strong link: If someone coughs then the conductor is annoyed

Result: The conductor is annoyed

We can now explain what is correct about the intuition that details should be omitted; *details should be omitted from the antecedent of the link*. And we can explain the apparently conflicting intuition that details should be included in explanations; *details should be included in the base*.

Notice that the result that good explanations are insensitive to initial conditions falls out of this account. The sensitivity of the explanation to the initial conditions depends on the antecedent of

the link. The more detailed the antecedent of the link, the more sensitive the explanation to initial conditions. This account recommends logically strong links, which means antecedents with fewer details, which means explanations which are insensitive to initial conditions.¹⁴

So far I have argued that logical strength is an explanatory virtue. Let's again ask the deeper question: why is logical strength an explanatory virtue? I suggest: logical strength is a fundamental virtue of *any* belief – it is good to be informed. Thus the *explanatory* virtue of informativeness is a special case of the doxastic (i.e. related to belief) virtue of informativeness.¹⁵

There is a further question: why think that logical strength is a fundamental virtue of beliefs? Here is a quick argument. It is very plausible that belief aims at truth.¹⁶ We can say that truth is a virtue of belief. (Metaethics uses 'norm' where philosophy of science uses 'virtue' i.e. truth is a norm of belief.) But there must also be other norms/virtues of belief – otherwise, the belief that $1+1=2$ would be as good as any belief could be (because it is true). What's lacking in the belief that $1+1=2$? It's not *informative*.¹⁷ It would be better to have a more informative belief (other things equal), so it is plausible that informativeness is a fundamental norm/virtue of belief. That is, logical strength is a fundamental norm/virtue of belief.

I'm not saying that logical strength (and truth) are the *only* fundamental norms/virtues of belief. Presumably evidence is also part of the story, to name just one more. But it is very plausible that logical strength is a virtue of belief. If so, the philosopher of science can rest content that a question about explanatory virtues has been reduced to a question about the norms of belief.

¹⁴ So the preference for logical strength also explains the preference for the counterfactual dependence relations emphasized by Woodward and Hitchcock (2003a, 2003b).

¹⁵ Compare Lewis 'What I have been saying [about explanatory virtues] applies just as well to acts of providing information about any large and complicated structure...The information provided, and the act of providing it, can be satisfactory or not in precisely the same ways.' (1986 p.227-8) Lewis is talking about pragmatics here, but I think the same applies to objective norms.

¹⁶ See Whiting 2012 for discussion and references.

¹⁷ Compare: 'Science does not aim, primarily, at high probabilities. It aims at a high informative content, well backed by experience'. (Popper 1954, p.146, original Italics.)

(In fact it might be better to say that logical strength is a norm of explanation than to say that logical strength is an explanatory virtue. It is plausible that explanatory virtues are features of hypotheses which provide a reason to believe the hypotheses i.e. which boost the probability of the hypotheses in some way.¹⁸ But being logically strong cannot be a reason to believe a hypothesis. If anything, being logically weak provides a reason to believe a hypothesis e.g. tautologies are maximally weak and should have credence of 1. So being logically strong cannot provide the kind of probability boost (to priors or posteriors) that explanatory virtues are usually taken to provide. So it's better to say that logical strength is a norm of explanation i.e. explanations ought to be strong. Still, the difference won't matter here so I will stick with the more familiar terminology of 'virtues'.)

Let's now work through the consequences for functional explanations.

4. Functionalism and Fodor's Argument

A *functional property* is a property that is characterized in terms of what it does. That is, functional properties are individuated by their causal profile.¹⁹

Fodor's (1997) argument for a preference for functional explanations is especially interesting:

[The] functionalist...still needs to explain why we should...prefer higher level closed laws (pain leads to avoidance) to lower level open laws (states that are $R_1 \vee R_2 \vee \dots$ lead to avoidance), all else equal. Why are we prepared to buy closed laws at the cost of reifying high level properties? My story is that this policy complies with an injunction that all of our inductive practice illustrates: Prefer the strongest claim compatible with the evidence, all else equal. p. 159

¹⁸ This is controversial; see Weisberg [2009] for discussion.

¹⁹ By 'functionalism' I mean the Ramsey-Carnap-Lewis theory described by Lewis (1970, 1972). This should be distinguished from other uses of 'functionalism' e.g. Shoemaker (2003), which I set aside.

I have been arguing for something like the claim that Fodor makes in the final sentence, though I would put it as follows: logical strength is a fundamental explanatory virtue.

But I want to argue that Fodor has got the implications for functionalism the wrong way round. Fodor suggests that functional links are logically stronger than lower level links. I will argue that functional links are logically *weaker* than lower level links.

In fact functional links are trivial, leading some to worry that they don't succeed as explanations at all. Ironically, Fodor himself had expressed this worry with a characteristically clear and colourful example that has become common currency²⁰:

The second problem is...a lurking suspicion that functional...explanation...is just too easy. Moliere had a finger on it:

In virtue of what did the morphine put Jones to sleep?...In virtue of its having dormative power. In what does having this dormative power consist? It consists in being the cause of a certain kind of effect; viz. in causing sleep. Fodor (1981) p.12

Italics original

The poor explanation of Jones being put to sleep is something like:

Basis: Jones took something with a dormative power

Link: Substances with dormative powers cause sleep (i.e. if someone takes a substance with a dormative power then it puts them to sleep)

Result: Jones was put to sleep

²⁰ Fodor goes on to give an answer by appealing to Turing machines, but this seems to restrict the scope of functionalism so much that it would not apply to any special sciences – with the possible exception of computing.

And the link tells us *nothing* about the world. It is analytic – a Carnap-sentence (see Carnap 1963 p.964, Lewis 1970, 1972).²¹ Thus functional links are logically weak, and Fodor’s affirmation that we should prefer logically stronger explanations ends up counting *against* functionalism. (And we can explain the intuitive unsatisfactoriness of the functional link in terms of its being uninformative – this fits with my position that informativeness is a fundamental explanatory virtue. But this doesn’t mean functional explanations are useless; see Bradley 2014)

Let’s investigate whether we could have an informative functional link. We could start with a link that mentions a realizer e.g.

Link: If something fires μ -opioid receptors then it puts people to sleep

And we could logically strengthen the explanation by adding a disjunct to the antecedent of the link:

Link: If something fires μ -opioid receptors or k-opioid receptors then it puts people to sleep

Don’t functional links have exactly this disjunctive structure, making them logically stronger than links without the disjunction?

No. Functional links do not have the structure above, in which adding a disjunct to the antecedent provides additional information about the realizers in the actual world. Functional properties are individuated by their causal powers, and are independent of the *actual* realizers.

We *can* use disjunctions to state something extensionally equivalent to a functional link if we refer to other possible worlds²²:

²¹ I do not intend to be committed to what is often called ‘analytic functionalism’ and is distinguished from ‘psycho-functionalism’ (see Levin 2016).

²² See Cohen 2002: ‘the dispositions a thing has depends crucially on the operative laws of nature: in worlds where the laws of nature are different, none of the [actual realizers] need serve as the basis for the disposition

Link: If something fires H receptors in humans in possible world 1 or fires J receptors in humans in possible world 2 or... then it puts them to sleep

This link is non-analytic but still necessary – it is true in all possible worlds, so it is maximally weak and uninformative.²³ So as far as the virtue of logical strength goes, it is just as bad an explanation as a functional explanation.

The views I wish to defend are now on the table. I have argued that logical strength is a fundamental explanatory virtue. This accounts for our intuitions in examples where omitting details looks virtuous – explanations can be improved by omitting details if (and only if²⁴) doing so makes the explanation logically stronger. The result is that higher level links formed by omitting details from lower level links provide better explanations; higher level links formed by using functional predicates provide worse explanations. It remains to generalize from these examples and respond to objections.

5. Two Generalizations

We need to make two generalizations to cover the examples discussed in the literature.

The *first* generalization extends our analysis to two other types of explanation. Like Hempel and Oppenheim, we focussed on *causal* explanations, which have antecedent conditions as bases and some kind of causal connection as links; but there are also *metaphysical* and *logical* explanations.

Metaphysical explanations have relatively fundamental states as bases and grounding principles as links. For example, the link ‘if there are molecules arranged table-wise then there is a

fragility. Consequently, we must take into account the (perhaps infinitely many) low-level configurations that serve as the basis for the disposition at each world $w \neq @$ as well’. p.80 (Cohen here assumes that the causal profiles of properties are not essential to them. Thanks to a referee for emphasizing this.)

²³ Fodor (1974) puts great stock in the distinction between infinite lists of realizers and functional properties, but the link is maximally weak either way. And this is another problem with the 1997 Fodor quote above. Fodor implies that higher level closed laws are logically stronger than lower level open laws; it seems to me that they must be extensionally equivalent, and so equally strong.

²⁴ I haven’t argued for this ‘only if’. There may be some other reason to omit details, though I don’t think there is.

table', combined with the base that molecules are arranged table-wise, explains that there is a table.

Logical explanations have minor premises as bases and material conditionals as links. For example, the assumption that 'if A then B', combined with the assumption that B, explains A.

	Basis	Link
Causal/nomological	Antecedent conditions	Law/generalization
Logical	Minor premise	Material conditional
Metaphysical	Relatively fundamental state	Grounding principle

Recall Putnam's original example:

Link: If it's a 1 inch square object then it does not fit through circular holes with diameter 1 inch

Although Putnam was defending higher level laws, he does not use a law of nature – it is not a law of nature that 1 inch squares don't fit through 1 inch holes. Nor is it merely a counterfactual-supporting generalization. It is much more like a grounding principle. To be more precise, it seems to follow from principles of geometry as applied to actual space, so perhaps we need to add mathematical principles to the list of causal, metaphysical and logical. The exact principles won't matter here though; what matters is that we need some kind of *conditional* in the explanation, and conditionals can be strengthened by weakening the antecedent. And this is why explanations that describe the overall shapes of the peg and board can be better than those that describe the positions of the molecules. Moving from causal to metaphysical or logical explanations doesn't change this.

The *second* generalization is needed because the concepts of the base can differ from the concepts of the antecedent of the link. Consider the explanation:

Basis: The card is scarlet

Link: If the card is red then the pigeon pecks

Result: The pigeon pecks

Logic does not connect scarletness with redness. We need *grounding principles*²⁵. These could be packed into the link, but it is clearer to separate them. We would end up with something like:

Basis: The card is scarlet

Grounding principle: Scarlet is a determinate of red

Link: If the card is red then the pigeon pecks

Result: The pigeon pecks

So the second generalization extends the account to explanations which require a grounding principle. The main point remains – omitting details from the antecedent of the link improves the explanation.

One might now argue that explanations are better without grounding principles. I won't argue that they are; but *if* they are, and assuming bases should contain lower level concepts (perhaps because they are more informative) it follows that links should contain lower level concepts. The result would be that explanations in terms of the lowest level science are best. This is the explanatory reductionist position that the best explanations are given in terms of the lowest level science. For example, assuming that redness comes in just two varieties – scarlet and maroon – the reductionist favours 'if the card is scarlet or maroon then the pigeon pecks' over 'if the card is red then the pigeon pecks'.²⁶

²⁵ See Schaffer (forthcoming).

²⁶ A different argument for this conclusion is given by Sider (2011), who argues that explanations should be given in joint-carving terms: "theories" based on bizarre, non-joint-carving classifications are unexplanatory' p.23 Assuming that lower level concepts carve at the joints better than higher level concepts, Sider's position supports explanatory reductionism. But I don't think non-joint-carving explanations are unexplanatory; I think they are *not as good* as joint-carving explanations.

What anti-reductionist arguments are there for preferring links that use higher level concepts? There are three main arguments for favouring higher level links: they avoid disjunctive predicates, illuminate generalities and are better at speaking to our interests (see Pereboom & Kornblith, 1991, Fodor 1974, 1997, Sober 1999, Clapp 2001, Batterman 2002²⁷).

These arguments all require further discussion, but none of them seem very strong. What's wrong with disjunctive predicates? Why is it better to use a single predicate that just ties a bow around the disjunction? And if the disjunction is heterogenous, is there really any generality to be illuminated, as opposed to an illusion of a generality created by a higher level concept? If there is a genuine generality, why can't that be illuminated by a lower level conditional connecting one effect with multiple causes? As for speaking to our interests, this is sensitive to the interests of the agent in question. It is hard to see that there are any morals to be drawn regarding objective explanatory virtues (I expand on objective vs. pragmatic virtues in the next section). But I will remain officially neutral on this debate about whether higher or lower level links are better; I claim only that logically strong links are better.

6. Should the Base Really be Maximally Strong?

These final three sections defend my account in more detail. This section defends the view that bases should be maximally strong, and the next (7) discusses Garfinkel's objections. The final section (8) defends the view that links should be maximally strong from the problem of disjunctive explanations.

Let me start by pointing to a possible means of retreat. My main thesis is that logical strength is a fundamental explanatory virtue, and this recommends adding information to any explanation. This thesis leaves open that there are other fundamental explanatory virtues, including those which recommend removing information from an explanation. So perhaps the correct level of information

²⁷ Batterman (2000, 2002, forthcoming) argues that the question the reductionist cannot answer is: (MR) How can systems that are heterogeneous at some (typically) micro-scale exhibit the same pattern of behavior at the macro-scale? This seems analogous to the one over many argument for Platonism: how can two different objects, a and b, have the same property F? I am inclined to give the nominalist answer: in virtue of a being F and b being F. Perhaps a similar nominalist style answer applies to Batterman's challenge.

depends on a trade-off between competing virtues. Having noted this path of retreat, I do not plan to take it. I am doubtful that there is any such competing explanatory virtue, so I will try to defend the thesis that explanations are always improved by adding information. This means that bases should contain as much information as possible, right down to the level of fundamental physics. This might appear odd. But I think it can be independently motivated.

Suppose an unusually curious Ancient Greek wants to know why Socrates died. She is told that he drank poison. Not fully satisfied, she asks which poison was drunk; a reasonable question, she is told he drank hemlock. But our unusually curious Ancient Greek is still not satisfied. What is the process by which hemlock killed Socrates? Surely another reasonable question, though not one any Ancient Greek could have answered. But let's pretend a modern biologist arrives and explains that hemlock contains coniine, which disrupts the workings of the central nervous system. Will our curious Ancient Greek be satisfied? Perhaps not. She now asks what the process is by which coniine disrupts the central nervous system. A chemist arrives to explain how coniine paralyzes muscles by blocking the nicotinic receptors which cause muscular contraction. And how does it do that, our curious Ancient Greek now asks. A physicist is now needed to explain the interaction between the coniine and the nicotinic receptors. So we have rapidly descended to physics to explain why Socrates died, and at each step the explanation has improved. At what point is our curious Ancient Greek satisfied? Perhaps only once we have descended to the level of fundamental laws, where we are faced with phenomena that cannot be further explained. And only then is the explanation *good* enough to satisfy the curious Ancient Greek.

Someone might object that this continual movement towards further information does not improve the explanation. After all, the behaviour of the curious Ancient Greek is not realistic; real people do not have either the time or the interest to pursue explanations into such details.

We can block this objection using two distinctions. First, distinguish an *ideal* explanatory text, which contains all the information that could be relevant to explaining the event, from the non-ideal explanations that we actually *give*, and which are explanatory in virtue of conveying information about

some portion of the ideal text.²⁸ Ideal explanatory texts are theoretical infinitely long sentences that cannot be made better or worse; we are discussing the virtue of a *given* explanation. Second, distinguish the virtues that make a given explanation *objectively better* from *pragmatic* features that make an explanation *useful*.²⁹

My account says that a *given* explanation *objectively* improves as it expands and moves closer to the ideal explanatory text; I allow that more information can *pragmatically* worsen a given explanation. For example, given explanations are subject to pragmatic principles of conversation.³⁰ Giving too much information flouts the second half of Grice's (1975 p.45) maxim of quantity: 'Do not make your contribution more informative than is required'. And what is required depends on what the subject wants to know.³¹ This accounts for the intuition that extra information can worsen the explanation.

Equally, a curious subject who wants to know more information will find a more detailed explanation more useful. Thus, Elliott Sober (1999), responding to Putnam's claim that micro-properties are not explanatory writes:

Perhaps the micro-details do not interest Putnam, but they may interest others, and for perfectly legitimate reasons. Explanations come with different levels of detail. When someone tells you more than you want to hear, this does not mean that what is said fails to be an explanation (p. 547).

²⁸ Railton 1981 p.240 Railton is not explicit about how much information is in the ideal text. I think it should contain every true sentence.

²⁹ For related comments about objective virtues, see Clapp 2001 p.135, Haug, 2011 p.1147; but also Pereboom & Kornblith, 1991 p.127 for dissent.

³⁰ Here we can apply the earlier Lewis quote and an extra sentence about pragmatics: 'What I have been saying [about explanatory virtues] applies just as well to acts of providing information about any large and complicated structure...The information provided, and the act of providing it, can be satisfactory or not in precisely the same ways. There is no special subject: pragmatics of explanation' (1986 p.227-8)

³¹ This seems to be Lewis's (1986a p.227) diagnosis of what's wrong with the micro-explanation of Putnam's peg.

I would add: when someone tells you more than you want to hear, the given explanation objectively *improves* as it gets closer to the ideal explanatory text.

The objector might continue³²:

‘Objection: Suppose we add to the explanation of Socrates’ death that he held the cup with his left hand. This addition makes the explanation more informative, but surely does not improve the explanation.’

In response, I maintain that the additional information does improve the explanation. I concede that this is a counter-intuitive result, but I think the counter-intuitiveness is due to the pragmatics and our limited cognitive capacity.

Perhaps the counter-intuitiveness can be undercut by consideration of cases where the additional information is needed for the explanation to succeed. Suppose you believe that anyone who drinks hemlock with their right hand is given the antidote by Dexter, a deity who loves the right-handed. Knowing that Socrates is right-handed, you are puzzled as to why the hemlock killed Socrates. Where was Dexter? Your puzzlement is resolved only when you have the additional information that Socrates drank with his left hand this time. Generally, given the right background beliefs, anything can explain anything – a kind of explanation holism. It follows that for any given piece of information and any phenomenon, there are states of background knowledge where that piece of information explains that phenomenon. So the more information in the base, the more states of background knowledge allow the explanation to succeed, and the objectively better the explanation. We have limited time and interest, so only want to hear information that is needed for the explanation to succeed given our background beliefs at the moment of explanation - but there is always more to know.³³

7. Anti-Reductionist Arguments Regarding the Base

³² Thanks to Ben Wolfson for pressing this point.

³³ Compare Lewis (1986a): ‘Your explanatory information is only partial. Yes. And so is any serving of explanatory information we will ever get . . . There is always more to know.’ p.237

My position is that the more information the base contains, the better (objectively) it is. Although he doesn't use the link/base distinction, Garfinkel (1981) gives one of the most detailed arguments in the literature for removing information from the base, so it is worth going through his reasoning in detail.³⁴ I will describe two arguments for removing information from the base, offer responses, then apply the responses to Garfinkel's example.

7.1. Lower level bases that omit information

Suppose we boil water in a sealed glass container until the glass cracks. Compare the following bases:

Basis 1: Molecule A hit the glass with momentum greater than z .

Basis 2: The water was boiling.

Basis 2 seems better. Let's grant that it is. One might argue that basis 2 is better because it has less information than basis 1. Counterexample?

No. *Basis 2 is not less informative than basis 1.* (Neither entails the other.) Basis 1's explanation in terms of molecule A omits any information about the *other* molecules – molecule A could have been moving freakishly fast by chance. So basis 1 is lower level but omits information. Basis 2, by saying the water was boiling, includes the information that the other molecules were moving fast as well. So this is not a case in which an explanation with less information is better. In fact neither is more informative than the other.

To devise a case where the higher level explanation that uses the concept of 'boiling' is less informative, the lower level explanation would have to provide the details of the trajectories of many molecules – at least enough to necessitate that the water is boiling. So we need to compare basis 2 with something like:

³⁴ If we interpret Garfinkel as arguing that we should omit details from the antecedent of the link then I agree with him.

Basis 3: Molecule A is at position L1 with momentum x_1 , Molecule B is at position L2 with momentum x_2 ...

Basis 3 provides at least as much information as basis 2, making the explanation at least as good on my account. And this seems intuitively correct – for example, explanations with basis 3 are just as counterfactually robust as explanations with basis 2. Extra details appear to make the explanation worse only when we make explicit *some* details (about molecule A) and fail to say anything about the *other* details (the rest of the molecules).

7.2. Contrast

One might think that explanations with extra details invoke the wrong contrast class. I will argue that the contrast class is separable from the details provided.

Let's grant that explanation is contrastive.³⁵ That is, p rather than q explains x rather than y . Recall the example where the pigeon pecks at a red card. Hold fixed that we want to explain why the pigeon pecked rather than didn't peck. We might try to contrast the scarlet card with a maroon card. That is (leaving the link implicit):

Failed microexplanation

Basis: The card is scarlet rather than maroon

Result: The pigeon pecked rather than didn't peck

This purported explanation fails. The pigeon would have pecked at any shade of red.

³⁵ Woodward's interventionist model of explanation is naturally thought of as contrastive. Schaffer 2005 defends contrastive causation. The difference between contrastive causation and contrastive explanation won't matter here. It also doesn't matter if explanation really is contrastive – I am conceding the point to the objector.

But it is a mistake to infer that any detailed explanation fails. We just have to make sure that the contrast class contains only possibilities where the pigeon does not peck e.g.:³⁶

Successful microexplanation

Basis: The card is scarlet rather than turquoise

Result: The pigeon pecked rather than didn't peck

7.3. Garfinkel's discussion

Let's now apply these two points to Garfinkel's discussion. Garfinkel's strategy is to argue for the failure of the microexplanation of why a rabbit was eaten. My reply is that he makes both the illicit moves above.

Garfinkel *first* mentions that the micro level contains a huge amount of information, and then explicitly omits most of it:

'the overall nature of the microlevel is a huge-dimensional determinism³⁷, which, given the complete description of all the equations of interaction between individual foxes and individual rabbits and given a complete specification of foxes and rabbits, tells us the individual destiny of every one of them at every future time. *Extracting from this mass* the data relevant to rabbit r, we learn that...the microexplanation is therefore something like:

Rabbit r was eaten because he passed through the capture space of fox f.' p.55

Italics added.

³⁶ See Franklin-Hall (2016 section 5).

³⁷ One might deny that the microlevel is deterministic, but it won't matter here.

By extracting from the mass of data, Garfinkel has omitted information from the lower level description.

Second, once Garfinkel has focussed on the rabbit's passing through the capture space of f, he invokes the contrast class of the rabbit passing through some other fox's capture space and concludes:

'The object of the macroexplanation is why the rabbit was eaten rather than not eaten, while all the microexplanation tells us is why the rabbit was eaten by fox f...rather than by some other fox' p. 56

So Garfinkel correctly rejects:

Failed microexplanation

Basis: Rabbit r passed through the capture space of fox f1 rather than fox f2

Result: Rabbit r was eaten rather than not eaten

In response, we can give a microexplanation that a) gives at least as much information as the macroexplanation, and b) has a contrast class (involving a small fox population) that explains why the rabbit was eaten rather than not eaten:

Successful microexplanation

Basis: Rabbit r was at location l10 and foxes f1...f200 were at l1...l200 respectively rather than rabbit r was at location l10 and foxes f1...f9 were at l1...l9 respectively.

Result: Rabbit r was eaten rather than not eaten

These points can be adapted to the other examples (though having conceded to the objector that explanation is contrastive, we have to modify some of the bases; we'll also leave the grounding principles of section 5 implicit, along with the links):

Peg:

Failed microexplanation

Basis: The peg was square with structure S1 rather than square with structure S2.

Result: The peg got stuck rather than going through the hole

Successful microexplanation

Basis: The peg was square with structure S1 rather than round with structure S3.

Result: The peg got stuck rather than going through the hole

Glass:

Failed microexplanation

Basis: Molecule A was the first³⁸ to hit the glass with momentum greater than z rather than molecule B was the first to hit the glass with momentum greater than z

Result: The glass cracked rather than stayed intact

Successful microexplanation

³⁸ This reference to the first molecule to hit the glass ensures that the contrasts are incompatible. (Molecule A hitting the glass with momentum greater than z is compatible with molecule B hitting the glass with momentum greater than z.) Thanks to x and y

Basis: Molecule A hit the glass with momentum greater than z rather than no molecule hit the glass with momentum greater than z

Result: The glass cracked rather than stayed intact

Conductor:

Failed microexplanation

Basis: Bob coughed and Alice didn't rather than Alice coughed and Bob didn't³⁹.

Result: The conductor is annoyed rather than serene

Successful microexplanation

Basis: Bob coughed rather than no-one coughed.

Result: The conductor is annoyed rather than serene

Thermodynamics:

Failed microexplanation

Basis: The gas had initial state A with velocities $V_1-V_{1,000,000,000}$ rather than initial state B with (different) velocities $V_1-V_{1,000,000,000}$.

Result: The gas filled the balloon rather than stayed in one spot

Successful microexplanation

Basis: The gas had initial state A with velocities $V_1-V_{1,000,000,000}$ rather than a billion stationary molecules.

³⁹ Why not just: Bob coughed rather than Alice coughed? Because these possibilities are compatible, so do not form a contrast.

Result: The gas filled the balloon rather than stayed in one spot

8. Should the Antecedent of the Link Really be Maximally Weak?

Compare the following explanations:

Good explanation

Base: Sandy is a man

Weak link: If Sandy is a man then Sandy does not get pregnant.

Result: Sandy does not get pregnant

'Problematic' explanation

Base: Sandy is a man

Disjunctive link: If Sandy is a man or Sandy takes birth control pills then Sandy does not get pregnant.

Result: Sandy does not get pregnant

According to my account, the disjunctive link is logically stronger than the weak link, so better than the weak link, and this looks counter-intuitive.

The most important response dialectically is that competing theories of explanation face a similar problem. After all, if omitting details is a virtue, then adding disjunctions should improve the explanation. So most theories of explanation struggle to explain where this Disjunctive Explanation goes wrong:

Disjunctive Explanation⁴⁰

⁴⁰ We should also distinguish this case from one where the disjunct is irrelevant:

Non-problematic explanation

Base: Sandy is a man or Sandy takes birth control pills

Disjunctive link: If Sandy is a man or Sandy takes birth control pills then Sandy does not get pregnant.

Result: Sandy does not get pregnant

Weslake (ms) calls this the problem of disjunctive explanations⁴¹ and shows it is a problem for Hempel and Oppenheim (1948), Kitcher (1981), Woodward (2003), Woodward and Hitchcock (2003b), and Strevens (2004; 2009).

My account does say something about why the Disjunctive Explanation is worse than the Good Explanation – the base is logically weak.⁴² So this example supports my theory over alternatives.

The problem case for me is where the base is held fixed and a disjunction is added to the link.

Repeating from above:

‘Problematic’ explanation

Base: Sandy is a man

Disjunctive link: If Sandy is a man or Sandy takes birth control pills then Sandy does not get pregnant.

Result: Sandy does not get pregnant

One possible response is to add a further explanatory virtue that rules out disjunctive links; we’ll consider a couple of ways to do that in parenthetical comments below. But I offer a different response:

Base: Sandy is a man

Strong link: If Sandy is a man or the Sun is yellow then Sandy does not get pregnant.

Result: Sandy does not get pregnant

I suggest we require conditionals in causal explanations to be stronger than the material conditional, which ensures this link is false, so fails as an explanation (assuming that only the truth can explain).

⁴¹ Actually Weslake’s example is slightly more complicated, but I take it the point is the same.

⁴² Objection: But this is balanced by the virtue that the link is logically strong. Reply: Yes, I’ve assumed the base is the trumping consideration. Perhaps there is a non-extreme exchange rate. And perhaps the exchange rate is context-sensitive. I set aside these complications.

I maintain that this disjunctive link *does* improve the explanation. Why might someone think that it doesn't? I can think of four reasons.

First, someone might object that the disjunctive link is not law-like, holding that explanations must be law-like. But the claim that explanations must be law-like is highly controversial. I'm inclined to a more liberal view – for example, Jenkins and Nolan (2008) give compelling arguments that there can even be backwards explanations. (Still, those who demand that links must be law-like can add this as a restriction and accept the rest of my account.)

Second, perhaps natural properties⁴³ in links explain better, and the disjunction 'Sandy is a man or Sandy takes birth control pills' is not natural. But I don't think the preference for natural links can be maintained. Many intuitively good explanations are non-natural. For example, functional explanations are non-natural⁴⁴. The explanation that someone died because they were poisoned is non-natural because so many different substances realize poisons.⁴⁵ Furthermore, some think it is costly to wheel in metaphysically substantive concepts like naturalness.⁴⁶ (Still, those who demand that links must be natural can add this as a restriction and accept the rest of my account.)

Third, we might be misled by focussing only on the link and ignoring the base. Franklin-Hall (2016) argues that interventionists are pushed towards disjunctive explanations, with the result that 'the interventionist is recommending very peculiar, and in some respects very uninformative, explanations, those that don't go any way towards specifying the particular circumstance that brought about the effect to be explained' (p.573). But on my account the link is not in the business of specifying the particular circumstances – that is for the base.

Fourth, dissatisfaction with disjunctive links may be due to the pragmatics. Recall that the ideal explanatory text contains every sentence that could be used in an explanation of the event.

⁴³ Lewis 1983

⁴⁴ My earlier objection to functional explanations is that they are logically weak, not that they are non-natural.

⁴⁵ Objection: Poisons are unified at the relevant (higher) level. What unifies poisons is their ability to disturb organisms. Response: But on such a liberal account of unification, being a man and taking birth controls are unified, as both prevent pregnancy.

⁴⁶ See Weslake (ms) for a detailed discussion.

When we ask for explanations, we will not be happy to be *given* the ideal explanatory text, because we don't have the time or cognitive capacity to handle it – we want the information that is relevant to us. The explanations above have the base 'Sandy is a man'. This naturally invokes a case where the explainer – the person doing the explaining – knows Sandy is a man. If this explainer gives us the disjunctive link, they have violated Grice's (1975 p.46) maxim: be relevant. They would violate the same maxim as someone who told you Bill was in the library or the pub, while knowing that Bill was at the pub.

But we can fill out the story so that the disjunctive link does not violate any pragmatic rules. Suppose an alien does not know whether men can get pregnant, doesn't know whether Sandy is a man or a woman, and is very interested in how pregnancy works, and why Sandy is not pregnant. They read a book on human biology. The first chapter tells them that men cannot get pregnant and women can. The second chapter tells them that women can avoid getting pregnant by taking birth control pills. When they learn that Sandy is a man, the explanation is complete. Would the explanation have been improved by omitting the second chapter of the book? Surely not. The second chapter gives the alien further knowledge about the causal dependencies associated with pregnancy; it gives the alien counterfactual knowledge about what would have happened if Sandy had turned out to be a woman taking birth control pills. So the disjunctive link is more useful to the alien than the weak link.

So there a number of reasons the disjunctive link might look problematic, but I think they can all be answered.

9. Conclusion

I have criticized the contemporary consensus that higher level explanations are better because they omit details. I agree that omitting details can be a virtue, but this is because omitting details can make an explanation logically stronger; logical strength is a fundamental explanatory virtue. So higher level explanations are better than lower level explanations when logically stronger and worse when logically weaker (other things equal).

One of the main arguments for anti-reductionism is the claim that omitting details makes an explanation better. If, as I have argued, it is logical strength that makes an explanation better, then this argument for anti-reductionism is undercut.

Acknowledgments

Thanks to Jonathan Schaffer and Robbie Williams for comments on a draft of this paper, to Roman Frigg, Laura Franklin-Hall, Michael Weisberg and Ben Wolfson, and to the participants at the Workshop on Theoretical Virtues at Leeds 2017, especially Lina Jansson, Juha Saatsi, Jonathan Tallant and Pekka Väyrynen. This work was funded by EU Marie Curie Re-integration grant #656441.

DARREN BRADLEY

School of Philosophy, Religion and History of Science

University of Leeds

Leeds LS2 9JT

UK

d.j.bradley@leeds.ac.uk

References

- Batterman, R. W. 2000a. "A 'Modern' (= Victorian?) Attitude towards Scientific Understanding." *Monist* 83 (2): 228–57.
- Batterman, R. (2000b). Multiple Realizability and Universality. *British Journal for the Philosophy of Science* 51 (1):115-145.
- Batterman, R. 2002. *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford: Oxford University Press.
- Batterman (2017). Autonomy of Theories: An Explanatory Problem. *Noûs* 51 (1).
- Carnap (1934) THE UNITY OF SCIENCE. Trans. with an introd. by M[ax] Black. [Psyche Miniatures,

General Series, no. 63.] London: Kegan Paul, Trench, Trubner: & Co.

Clapp, L. (2001). Disjunctive properties: Multiple realizations. *The Journal of Philosophy*, 98(3), 111-136.

Clarke, C. (2016) The explanatory virtue of abstracting away from idiosyncratic and messy detail. *Philosophical Studies* 173:6, 1429-1449.

Craver, C. F. (2006), 'When mechanistic models explain', *Synthese*, 153, 355-76.

--- (2007), *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience* (Oxford: Clarendon Press).

Darden, L. & Craver, C. (2002). Strategies in the Interfield Discovery of the Mechanism of Protein Synthesis. *Studies in History and Philosophy of Science Part C* 33 (1):1-28

Darden, L. (2008), 'Thinking Again About Biological Mechanisms', *Philosophy of Science*, 75, 958-69.

Dizadji-Bahmani, F., Frigg, R. & Hartmann, S. (2011). Confirmation and reduction: a Bayesian account. *Synthese* 179 (2):321-338.

Fodor, J. A. 1974. "Special Sciences; or, The Disunity of Science as a Working Hypothesis." *Synthese* 28 (2): 97–115.

———. 1997. "Special Sciences: Still Autonomous after All These Years." *Nous* 31:149–63.

Franklin-Hall, L. R. (2016). High-Level Explanation and the Interventionist's 'Variables Problem'. *British Journal for the Philosophy of Science* 67 (2):553-577.

Garfinkel, A. 1981. *Forms of Explanation: Rethinking the Questions in Social Theory*. New Haven, CT: Yale University Press

Grice, H.P. (1975). Logic and conversation. In: P. Cole and J.L. Morgan (eds.), *Syntax and Semantics 3: Speech Acts*, 41-58. New York: Academic Press.

Haug, M. C. (2011). Abstraction and Explanatory Relevance; or, Why Do the Special Sciences Exist? *Philosophy of Science*, 78(5), 1143–1155. <https://doi.org/10.1086/662257>

Hempel, C. G., and P. Oppenheim. 1948. "Studies in the Logic of Explanation." *Philosophy of Science* 15 (2): 135–75

- Jackson, F., and P. Pettit. 1992. "In Defense of Explanatory Ecumenism." *Economics and Philosophy* 8 (1): 1–21.
- Jenkins, C. S., & Nolan, D. (2008). Backwards explanation. *Philosophical Studies*, 140(1), 103-115.
- Kim, J. (1992) "Multiple realization and the metaphysics of reduction," *Philosophy and Phenomenological Research* 52, pp. 1-26.
- Kim, J. (1993), *Supervenience and Mind*, Cambridge Studies in Philosophy, Cambridge University Press, England.
- Kim, J. 1998. *Mind in a Physical World*. Cambridge, MA: MIT Press
- Kitcher, Philip. 1984. "1953 and All That. A Tale of Two Sciences." *Philosophical Review* 93:335–373.
- Levy, Arnon & Bechtel, William (2013). 'Abstraction and the Organization of Mechanisms'. *Philosophy of Science* 80 (2):241-261.
- Lewis, D. (1986a) "Causal Explanation," in Lewis 1986b, pp. 214–240.
- Lewis, D. (1986b) *Philosophical Papers, Volume II*, Oxford: Oxford University Press
- MacDonald, G. 1992. "Reduction and Evolutionary Biology." In *Reduction, Explanation and Realism*, edited by David Charles and Kathleen Lennon, 69–96. Oxford: Oxford University Press.
- Machamer, Peter, Darden, Lindley, and Craver, Carl F. (2000), 'Thinking About Mechanisms', *Philosophy of Science*, 67, 1-25.
- Nagel, E., 1951, *The Structure of Science*, New York: Harcourt, Brace and World.
- Nozick, Robert 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Oppenheim, P. and H. Putnam, 1958, "The unity of science as a working hypothesis", in H. Feigl et al. (eds.), *Minnesota Studies in the Philosophy of Science*, vol. 2, Minneapolis: Minnesota University Press
- Pereboom, D., & Kornblith, H. (1991). The Metaphysics of Irreducibility. *Philosophical Studies*, 63(August), 125–45.
- Putnam, Hilary (1967), "Psychological Predicates", in W. Capitan and D. Merrill (eds.), *Art, Mind, and Religion*. Pittsburgh: University of Pittsburgh Press, 37-48. Reprinted as "The Nature of Mental

- States" in *Mind, Language, and Reality*. Cambridge: Cambridge University Press, 1975, 429-440.
- Railton, P., 1981, 'Probability, Explanation, and Information.', *Synthese*, 48: 233–56
- Sawyer, R. Keith. 2002. "Nonreductive Individualism: Part I—Supervenience and Wild Disjunction." *Philosophy of the Social Sciences* 32:537–559.
- Schaffer, J. (forthcoming) *The Ground Between the Gaps*, *Philosophers' Imprint*
- Skyrms, Brian 1980. *Causal Necessity: A Pragmatic Investigation of the Necessity of Laws*. New Haven, CT: Yale University Press.
- Shoemaker, S., 2003 "Realization and Mental Causation", in *Identity, Cause, and Mind*, Expanded Edition, Oxford: Clarendon Press.
- Sober, E. 1999. "The Multiple Realizability Argument against Reductionism." *Philosophy of Science* 66 (4): 542–64.
- Strevens, M. 2000. "Do Large Probabilities Explain Better?" *Philosophy of Science* 67 (3): 366–90.
- . 2004. "The Causal and Unification Approaches to Explanation Unified—Causally." *Nous* 38 (1): 154–76.
- . 2009. *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press
- . 2011 'Depth' *Montreal Review* <http://www.themontrealreview.com/2009/Depth-an-account-of-scientific-explanation-Michael-Strevens.php>
- . 2014 'Probabilistic Explanation' in *Physical Theory: Method and Interpretation* Sklar (ed.), Oxford University Press
- Weisberg, J., 2009. "Locating IBE in the Bayesian Framework," *Synthese*, 167: 125–143
- Weslake, B. (2010). Explanatory depth. *Philosophy of Science*, 77(2), 273-294. Chicago
- (ms) 'The Problem of Disjunctive Explanations'
- White, R. (2005). Explanation as a Guide to Induction. *Philosophers' Imprint*, 5(2), 1-29.
- WHITING, D. (2012), DOES BELIEF AIM (ONLY) AT THE TRUTH?. *Pacific Philosophical Quarterly*, 93: 279–300. doi:10.1111/j.1468-0114.2012.01421.x

- Woodward, J. 2003. *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.
- . 2008. "Mental Causation and Neural Mechanisms." In *Being Reduced: New Essays on Reduction, Explanation, and Causation*, ed. J. Kallestrup and J. Hohwy, 218–62. Oxford: Oxford University Press.
- Woodward, J., and C. Hitchcock. 2003a. "Explanatory Generalizations," pt. 1, "A Counterfactual Account." *Nous* 37 (1): 1–24.
- . 2003b. "Explanatory Generalizations," pt. 2, "Plumbing Explanatory Depth." *Nous* 37 (2): 181–99.
- Yablo, Stephen 1992. Mental Causation, *The Philosophical Review* 101/2: 245–80.
- Ylikoski, P. and Kuorikoski, J. 2010 'Dissecting explanatory power' *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, Vol. 148, No. 2 pp. 201-219