

LEVELLING COUNTERFACTUAL SCEPTICISM

Alexander Sandgren and Katie Steele

forthcoming in Synthese

In this paper, we develop a novel response to counterfactual scepticism, the thesis that most ordinary counterfactual claims are false. In the process we aim to shed light on the relationship between debates in the philosophy of science and debates concerning the semantics and pragmatics of counterfactuals. We argue that science is concerned with many domains of inquiry, each with its own characteristic entities and regularities; moreover, statements of scientific law often include an implicit ceteris paribus clause that restricts the scope of the associated regularity to circumstances that are 'fitting' to the domain in question. This observation reveals a way of responding to scepticism while, at the same time, doing justice both to the role of counterfactuals in science and to the complexities inherent in ordinary counterfactual discourse and reasoning.

Counterfactual scepticism, the thesis that most counterfactuals are false, has received a fair amount of attention recently. The discussion of counterfactual scepticism and how one should react to it, either by resisting it somehow or by accepting it and living with the consequences, yields insights into the nature and utility of counterfactuals and counterfactual reasoning.

In this paper, we present a response to counterfactual scepticism. Our proposal is motivated and informed by considerations from the philosophy of science, as well as by considerations primarily discussed in the literature on the semantics and pragmatics of counterfactuals. By bringing together insights from these two literatures we arrive at an attractive response to counterfactual scepticism that sheds light on both the semantics of counterfactuals and the role counterfactual reasoning plays in scientific inquiry.

1. The Road to Counterfactual Scepticism

One prominent defender of counterfactual scepticism is Alan Hájek. The following line of argument is substantially Hájek's (ms). Consider 1a:

(1a) Were you to drop that glass, it would shatter.

Ordinary speakers are inclined to think that 1a is true (at least on the understanding that the glass is fragile, and it would drop from a sufficient height onto a hard floor). Or consider another example from biology (concerning sexually reproducing species that are genetically diploid):

(2a) Were two individuals of some species to mate, the offspring would receive only one allele per trait from each parent.

2a invokes Mendel's Law of Segregation. We (and biologists for that matter) are inclined to accept 2a as true.

Or consider 3a uttered by a group of friends who were sorry that Alice was not at a party:

(3a) Were Alice to come to the party, it would be fun.

Many are inclined to accept 3a (on the understanding that Alice is generally very personable and fun-loving).¹

Nevertheless, many people, at least when pressed, tend to accept that in the counterfactual scenarios described above, alternative outcomes to those stated might have occurred. For instance, the laws of statistical mechanics allow that there is some possible microstate corresponding to the glass being dropped such that it takes a different trajectory and lands on a soft surface. Even under conditions conducive to ordinarily successful mating, certain biological irregularities may be such that an organism inherits something other than one allele from each parent for some trait. In the

¹ The glass example is a variant of an example discussed by Hájek (ms). The biology example appears in discussions of laws in the special sciences (see Rosenberg & McShea 2008). The example involving Alice's attendance at the party is much discussed in the literature on indicative and counterfactual conditionals.

case of the party, there might have been a house fire or a plumbing problem, such that the party, even with Alice's presence, was stressful rather than fun.

If alternative outcomes *might* have occurred, then the outcome described in the consequent *might not* have occurred. That is, it follows from the above observation that many, on pain of inconsistency, accept the following counterfactuals as true:

(1b) Were you to drop that glass, it *might not* shatter.

(2b) Were two individuals of some species to successfully mate, the offspring *might not* receive only one allele per trait from each parent.

(3b) Were Alice to come to the party, it *might not* have been fun.

What is more, the truth of a might-not statement seems to imply that the corresponding would-counterfactual is false. Yet, without prompting, speakers tend to utter these would-counterfactuals. That is the tension: why affirm the would-counterfactual if one would concede the truth of the apparently contradictory might-not claim?

Hájek upholds the might-not-counterfactuals in the face of this tension. He suggests that whenever a might-not-counterfactual is true, the corresponding would-counterfactual is false. Finally, Hájek argues that these cases in which would-counterfactuals are undermined by their might-not counterparts are *ubiquitous*.

Consider a simple case involving a coin toss.

(4a) Were you to toss that coin, it *would* come up heads.

(4b) Were you to toss that coin, it *might not* come up heads.

Since coin tossing is recognised as chancy, 4a rings false (the coin's landing heads is just *one* of the possible outcomes) while 4b rings true (it might land tails!). Hájek then notes that although the

coin-toss case invites the idea that the coin is fair, the same argument goes through even if the coin is heavily biased; even if the probability of its landing heads if tossed is .99, its landing tails is still a genuine possibility, so 4a is false and 4b is true.

An overwhelming reason to think that a process is chancy would be if its dynamics are genuinely *indeterministic*. A system is indeterministic just in case the state of the system at a time plus the laws of nature do not determine subsequent states of the system. But chanciness can also arise in a deterministic setting. In the coin-toss case, we might think of the chanciness as arising from treating the antecedent as picking out a *class* of tosses, where these tosses differ along parameters like momentum, velocity, etc. Hájek describes such antecedents as *unspecific*. Understood in this way, the coin case is, again, chancy; some ways of tossing the coin lead to its landing heads and others lead to its landing tails.

Hájek argues that once we recognise this about coin tosses we should recognise that nearly every counterfactual scenario is similarly chancy, either because the relevant process is not deterministic, or because the relevant antecedents are unspecific, in Hájek's sense. So for almost all counterfactual scenarios, there are multiple possible outcomes, each of which might not occur. This, combined with the claim that whenever a might-not-counterfactual is true the corresponding would-counterfactual is false, leads Hájek to claim that most counterfactuals are false.²

2. Our Proposal in Outline

In this section, we outline our response to counterfactual scepticism. In subsequent sections, we spell out the proposal in more detail, compare it to rival responses, and point to some of its important features.

² Why *most* counterfactuals rather than *all* counterfactuals? According to Hájek, there are a number of exceptions including counterfactuals with necessary consequents (e.g. if snow were orange, 2 and 2 would make 4), counterfactuals with impossible antecedents (e.g. were I to square the circle, I would be a poached egg) and counterfactuals with explicitly probabilistic consequents (e.g. were you to drop that glass, it would probably shatter).

We propose that counterfactuals should be interpreted relative to a *domain* of scientific inquiry, where this domain marks out both a *level* or grain of scientific inquiry and a *scope* for the inquiry. We use the term ‘scientific’ here rather liberally. Some of the domains we have in mind are well-recognised as scientific, while others are not. What is important is that the domains mark out recognisable regularities or patterns in the world, be they of greater or lesser significance.³ One highly significant domain is fundamental physics, taken to be the finest grain of inquiry and of universal scope. Another domain is human psychology, where the grain and scope concern human minds. A narrower domain still is the dynamics of groups involving Alice. Indeed, there are a multitude of relatively insignificant ‘homely’ scientific domains, such as my sleeping patterns, or your behaviour in response to alcohol, or the promptness of service at a particular café. The more established and significant the scientific domain, the more its level and scope are settled and well recognised, at least by experts. For the more homely domains, there may be considerable ambiguity about the content.

Depending on the domain of inquiry, the antecedent of a counterfactual statement is interpreted differently: it is subject to an implicit ‘fitting circumstances’ clause tied to that domain. For instance, we propose that the counterfactual statement 1a be interpreted as expressing a proposition of the form 1s.

(1s) Were you to drop that glass *under fitting circumstances for domain d* , it would shatter.

In this case, given the terms used (‘glass’ and ‘shatter’) in the counterfactual utterance, the domain in question is arguably ‘macro-object dynamics’, which we can denote d_m . That is to say that 1a expresses 1s, where the domain d is d_m .

On one plausible reading (which we elaborate in the next section), the *fitting circumstances* amount, roughly, to ideal conditions in which the regularity or regularities pertinent to the domain

³ Partly for that reason we will drop the term ‘scientific’ from here on in, and use the simpler phrase ‘domain of inquiry’, or just ‘domain’.

in question hold in isolation, absent *interfering* factors. When we are interested in the physical properties and mechanics of everyday objects like fragile glasses, for instance, certain physical microstates that interrupt macro-object dynamics count as *interferences*; these events are excluded in the antecedent, and do not therefore confound the relevant regularity. So 1a, interpreted as above, is true.

If 1a were expressed in a different setting, within the context of a statistical mechanics lecture, for instance, the pertinent domain may rather be statistical mechanics, denoted d_s . In that case 1a expresses 1s, where the domain d is d_s . Statistical mechanics concerns all physically possible microstates; as such, no particular physically possible microstate would count as an interference and be excluded from the set of fitting circumstances. So 1a, interpreted in this way, is false.⁴

If we are right, interpreting a counterfactual requires identifying a domain of inquiry, which will make a difference to the content of the antecedent. This means that the *truth conditions* of counterfactual claims depend on the pertinent domain of inquiry. In turn, the *truth values* of counterfactuals depend on what *lawful regularities* actually hold. Our proposal is a species of contextualism. This will be even more apparent below. We will return to the difference between our proposal and other forms of contextualism in section 5.

One might anticipate how our proposal yields a response to counterfactual scepticism. We saw that on the most natural reading of 1a it is true, because the fitting circumstances do not permit any strange trajectory for the vase. A similar story can be told for 2a and 3a. In the former case, the relevant domain of inquiry plausibly involves biological individuals subject to the lawful regularities, involving sexual reproduction, inheritance, etc., of ordinary macro-biology. As such, circumstances involving aberrant recombination of alleles, despite being well understood and of interest within the domain of chemistry, or even microbiology, are not fitting. In the case of 3a, the relevant domain of inquiry is plausibly one concerning Alice's impact on others in ordinary social milieux. Here, circumstances in which the party is ruined by a house fire or a plumbing

⁴ At the most fundamental level of inquiry the fitting circumstances clause may be vacuous in that no possible event or sequence of events counts as an interference.

problem are arguably not fitting, despite, perhaps, being relatively common events (although see below for further subtleties regarding this case). This serves to highlight one point that will come up again in the next section: fitting circumstances are not primarily identified by their high probability (although it may be that a set of circumstances can only count as fitting if it has sufficiently high probability).

Part of the puzzle spelled out earlier is that it often seems reasonable for speakers to reconsider a would-counterfactual's truth when it is pointed out that there are relevant possibilities in which the antecedent holds and the consequent does not. The contextualist's line is that there is a subtle shift in context in these cases such that apparently contradictory statements are both true.⁵ On our specific proposal, what is going on is that the interpreter is prompted to entertain a subtle shift in domain.⁶ Consider again the dropped-glass case. If one were to shift one's domain of inquiry from macro-object dynamics to statistical mechanics (as one might be prompted to do to match the perspective of a persuasive interlocutor), the set-up involving the glass is perceived differently, or in other words, the fitting circumstances differ: less-familiar possibilities for the trajectory of the glass may no longer count as interferences. On this reading then, 1a is true, or so we suggested above.

A general challenge for contextualist accounts (of everything from knowledge ascriptions to generics) is specifying how to identify and discriminate between contexts. For our account, the crucial question is: What domain of inquiry is appropriate for interpreting a counterfactual utterance? We cannot hope to answer that question in full here. Our analyses of the counterfactuals in this paper must therefore be taken as tentative. The aim is merely to provide an account that has

⁵ See K. Lewis (2016) for an eloquent recent proposal to this effect, which we discuss further in section 5. Note that K. Lewis (2016, especially section 4) provides an extended defence of the contextualist reading of sequences like 1a and 1b in the face of varied linguistic data, some of which suggests that the relevant statements should *not* both be taken as true.

⁶ However, given the nature of domains, we predict that speakers may sometimes be more reluctant to concede the falsity of an asserted would-counterfactual than is generally appreciated. This point is elaborated in section 5.

the basic resources for resolving (or otherwise dissolving) well-known puzzles concerning counterfactuals. We indicate how the story about domains might go, but it would be a much more substantial project to fully flesh out the details.

Furthermore, while we often exercise a certain kind of charity when discussing examples, we think it is somewhat open how much charity is appropriate. In any case, our account is not so flexible that all uttered counterfactuals come out true. For starters, there will be cases where the implied domain of inquiry and the laws associated with that domain are perfectly clear, and yet the antecedent, interpreted as we suggest, does not entail the consequent. It may even be that the consequent is not possible given the antecedent and relevant laws. Or else, the consequent may be possible, but not certain to occur given the antecedent, due to the antecedent and relevant laws describing a chancy process. For example, 4s, the interpretation of the statement 4a we propose, is false:

(4s) Were you to toss that coin *under fitting circumstances for domain d*, it *would* come up heads.

The coin *might not* land heads if *d* is the domain capturing the chancy behaviour of ordinary coin tosses, so 4s is false; that is, even in the absence of interferences (such as the coin vaporising in mid-air), the coin *might not* land heads.

In other cases, it may not be entirely clear which domain a counterfactual should be interpreted with respect to, making its truth conditions harder to assess. Sometimes confusion stems from the inherent imprecision of conversation; the speaker may be interpreted as positing quite different claims, depending on what is taken to be the domain of inquiry, where only one (or some) of these interpreted counterfactuals are true. In such cases what the crucial domain is can often be clarified by further conversation.

Sometimes the ambiguity about the relevant domain is of a deeper kind. There may be some domain of inquiry at play that is nonetheless unclear. Perhaps it is unclear *what* domain of inquiry is appropriate to the conversation. Or else the appropriate domain has vague boundaries and thus

its associated laws are unclear. For instance, this is arguably the case for 3a. On our proposal, this counterfactual is interpreted as follows:

- (3s) Were Alice to come to the party *under fitting circumstances for domain d* , it would be fun.

There is plausibly some regularity in the world concerning Alice's social impact at parties that is being invoked here and is the best candidate for domain d . But what precisely is this domain d ? Or rather: how should the regularity in question, or the reference class marked out by the fitting circumstances, be delineated? Do the fitting circumstances include only well-organised, fortuitous parties where nothing goes amiss, or do they include also badly organised parties where unforeseen predicaments like house fires occur? This will make a difference to the truth conditions of 3a. We think it appropriate to be permissive in evaluating such counterfactuals, in the following sense: if there is a respectable reading on which they are true, then they are not instances that provide clear support for counterfactual scepticism.⁷

In summary, we propose that a statement like 'if A were true then C would be true' must be interpreted with respect to a domain of inquiry d , and accordingly represented as ' $A_d \Box \rightarrow C$ ', where A_d is an interpretation of the antecedent as involving a fitting circumstances clause tied to domain d , and where the *box arrow*, $\Box \rightarrow$, represents the counterfactual connective.

Since our proposal concerns the interpretation of the antecedent and not the interpretation, as it were, of the box arrow, it can be plugged into a wide range of accounts of the semantics of counterfactuals. For instance, our proposal can be used to augment the strict-conditional analysis of counterfactuals according to which $P \Box \rightarrow Q$ is true just in case all P worlds are Q worlds. Accordingly, $P \Box \rightarrow Q$ is true just in case $\Box(P \supset Q)$ is true. Augmenting this view in light of our

⁷ These cases of 'deep' ambiguity or vagueness call for a more thorough examination that is beyond the scope of this paper. There are a number of subtly different stances one could take here depending on i) one's commitments regarding the nature of scientific domains and associated laws, and ii) one's preferred approach to vagueness.

proposal involves understanding a statement of the form ‘if A were true then C would be true’, interpreted with respect to domain of inquiry d , as equivalent to $\Box(A_d \supset C)$. Our proposal can also serve to augment other accounts of counterfactuals like the popular closest-worlds account of counterfactuals drawn from the work of Lewis (1973) and Stalnaker (1968), Kratzer-style (2012) modal semantics for counterfactuals, or causal-modelling-based semantics recently discussed by Briggs (2012) and Santorio (2019).

We are also assuming the usual characterization of might-not-counterfactuals, as the duals of would-counterfactuals. That is, $P \Diamond \rightarrow \neg Q$ entails $\neg(P \Box \rightarrow Q)$. On our proposal, the statement ‘If A were true, C might not be true’, interpreted with respect to domain d , is represented as $A_d \Diamond \rightarrow \neg C$. This entails that the statement ‘If A were true then C would be true’, also interpreted with respect to domain d , is false; that is $\neg(A_d \Box \rightarrow C)$.

3. Counterfactuals vis-à-vis Scientific Claims

We can now resist the sceptic’s conclusion. But one might worry, in light of Hájek’s argument, that we gave up something important to avoid scepticism. In particular, one might worry that our analysis of counterfactuals is less *objective* than is required, since, on our proposal, the content of counterfactual statements is tied to domains of inquiry.

Hájek argues that support from science is an important desideratum for an account of counterfactuals. The idea is well expressed by Kutach (2005): ‘What makes counterfactuals especially suitable for science is that the truth of counterfactuals depends largely on the general patterns that science aims to describe.’⁸ Others go further in claiming that a *defining feature* of the general patterns that science aims to describe – what makes these patterns or regularities lawful – is precisely their stability across counterfactual situations (see, e.g., Lange 2000). In any case, there is plausibly a tight connection between counterfactual dependence and the regularities science aims to describe.⁹ So far we agree with Hájek.

⁸ This quote appears in Hájek (ms).

⁹ For these naturalistic reasons, we will not address proposals in the literature that appeal to non-

We part ways concerning the implications of this general approach. According to Hájek, the ‘widespread falsehood [of would-counterfactual claims] is what one gets in a chancy world, as science teaches us that ours is’. But we contend that science does not teach us that *all lawful regularities are chancy*, once it is recognised that there are multiple autonomous domains of inquiry that differ in level or scope. In what follows we support this position by appeal to views (more and less widely held) within the philosophy of science literature. Our denial of the ‘widespread falsehood’ of would-counterfactual claims admittedly rests on an extension of this picture of science to include domains that are not typically recognised as scientific. We will not directly argue for the extension. Our strategy in this section is to outline the most compelling accounts of multiple autonomous domains of science, or *special sciences*, such that the reader sees the extension to the more homely domains, or *extra-special sciences*, as natural, or at least plausible.

At a superficial level, no-one disputes there being multiple domains of science. We recognise differences in the *level* of scientific inquiry. For example, the organisms studied in biology are clearly coarser-grained than the molecules studied in chemistry. We also recognise differences in the *scope* of scientific inquiry. For instance, it is not clear which of island biogeography and rainforest ecology is the more fine-grained inquiry, but the scope or subject matter of these two scientific domains differ. That said, no-one supposes that one could delineate and enumerate all the various scientific domains. Philosophers of science and scientists alike freely talk of domains of inquiry even if these domains have somewhat vague boundaries.

What is controversial is the status of ‘laws’ at domains other than fundamental physics. It is important for our account of counterfactuals that the higher domains, or special sciences, at least concern *lawful* regularities, i.e., non-accidental regularities that have some autonomy (whatever other conditions they may need to satisfy to count as full-blown laws). Some philosophers of

natural primitive metaphysical facts about what would have happened in various counterfactual circumstances, such as the proposal recently defended by Stefánsson (2018). See also Hawthorne (2005), Schulz (2010), and Moss (2013).

science have indeed defended this position, typically by appealing to *ceteris paribus laws* (*cp laws*), so-named because they have an implicit scope restriction to circumstances in which the regularity in question holds. Positive accounts of *cp laws* are many and varied, however (for a summary, see Reutlinger et al. 2017), so the burden falls on us to articulate and motivate an account of *cp laws* that supports our proposal. In what follows we take up this challenge. We identify a general way of conceiving *cp laws*, (as well as some possible ways to fill in the details) that is independently attractive (3.1) and supports our proposed response to counterfactual scepticism (3.2).

3.1 The content of *ceteris paribus* clauses

Consider some paradigmatic examples of *cp laws*:

Mendel's Law of Segregation: ‘In a parent, the alleles for each character separate in the production of gametes, so that only one is transmitted to each individual in the next generation’ (Rosenberg & McShea 2008, 36)

The Law of Demand: ‘Under the condition of perfect competition, an increase of demand of a commodity leads to an increase of price, given that the quantity of the supply of the commodity remains constant’ (Roberts 2004, 159; Kincaid 2004, 177)¹⁰

The Character of Shield Volcanoes: ‘Shield volcanoes erupt effusively’ (Strevens 2014, 1819)

The hidden *ceteris paribus* clause in each case restricts the circumstances at issue. The possibility of other circumstances in which the regularity is not manifest does not necessarily threaten the truth of the law. Mendel’s law of segregation does not apply, for instance, when chromosomes fail to properly separate. The law of demand does not apply in cases where consumers are *irrational*.

¹⁰ The *Law of Demand* and *Mendel’s Law of Segregation* both appear, with the references as stated, in Reutlinger et al. (2017).

Shield volcanoes with pyroclastic shields erupt explosively rather than effusively.¹¹ Nonetheless, these *cp* laws arguably say something substantive and interesting about how the world is.

But what exactly is the content of a *cp* clause? If the regularities that are picked out by *cp* laws are to play a role in determining the truth of counterfactuals, their content had better stand up to scrutiny, such that they indeed amount to non-accidental regularities. A basic worry along these lines turns on whether the *cp* clause should be read as *definite* or *non-definite*. A *definite cp* clause enumerates circumstances (typically, those excluded) in a concrete way while a *non-definite cp* clause merely describes the (excluded) circumstances in vague, open-ended terms. Either way, it is assumed that the *cp* clause describes a set of circumstances that amount to necessary and sufficient conditions for when the stated regularity does (not) hold. The supposed dilemma is as follows: there are no true definite *cp* laws, because some conditions in which the relationship does not hold will inevitably be overlooked, while non-definite *cp* laws are trivial, because they effectively have the form ‘*X* holds unless it doesn’t’. So *cp* ‘laws’ are either trivial or false – either way they are not genuinely lawful.

There have been different responses to this supposed dilemma, informed by subtly different accounts of *cp* laws (Reutlinger et al. 2017). Here we describe what we take to be the most promising general response, common to leading accounts of *cp* laws. The response involves reconceiving the *ceteris paribus* clause as implicitly referring to a fitting set of circumstances in which the purported regularity *does* in fact hold. So the *cp* clause identifies merely sufficient, as opposed to necessary *and* sufficient conditions, for the regularity to hold. This reconception of the *cp* clause, being less ambitious, is not subject to the dilemma described above.

But what would make a set of circumstances special or *fitting*, such that, if a regularity were to hold across this set, it may be considered lawful? This is a significant challenge. Fodor (1991) exploits the fact that *cp* laws pertain to ‘higher-levels’ of inquiry than fundamental physics, and as such, involve more or less coarse-grained entities that are *multiply-realised* at the finer grains of

¹¹ When a shield volcano is ‘pyroclastic’ it tends to erupt *explosively*, since it is not made by flows of basaltic lava with low gaseous content, as is more common among shield volcanoes.

analysis. He claims that, for the relationship ‘whenever A then B ’ to count as a *ceteris paribus* law, it must be the case that, for all possible instantiations or realisers of A , namely, $A(R_i)$, there *exists* circumstances C_i (a ‘completer for $A(R_i)$ ’) such that $A(R_i) \& C_i$ entails B . Admittedly this account of the *cp* clause is rather eccentric. The set of circumstances given by $A(R_i) \& C_i$ certainly has a special character, and ‘whenever A then B ’ comes out true under these circumstances. But the *cp* laws on this account will not always correspond to what are ordinarily considered lawlike regularities.

We regard the *dispositional* account as more promising, not least in better according with what are ordinarily considered lawlike regularities. On this account, the *cp* clause specifies circumstances under which the disposition or force underpinning the regularity is the only one that is manifested; there are no *interfering factors* that might otherwise compromise the regularity. For instance, ‘*cp*, the velocity of a falling object after elapsed number of seconds t is (approximately) 9.8 metres per second per second multiplied by t seconds’ says that in the special circumstances whereby gravity acts in isolation (for starters, when air resistance is nil), objects at the specified time have the stated falling velocity. The *cp* clause can in this way be understood as a statement of merely sufficient conditions in which the regularity in question holds. Moreover, as per Fodor’s account, the fact that the regularity in question holds in these circumstances may be regarded as conferring lawful status on the regularity, due to the special representative status of the circumstances. (On the dispositional account, there is supposedly also a sense in which *cp* laws *always* hold for the entities involved – in the form of a potentiality – which further bolsters the case for their being genuine laws.)

The *completer* and *dispositional* accounts of *cp* laws are different in many ways, but they both get something right – the implicit *cp* clause featuring in *cp* laws identifies special sufficient conditions, or *fitting circumstances*, in which the purported relationship holds. These accounts offer alternative ways, distinct from high probability, that a set of circumstances may be fitting. That is not to say that we rule out other notions of fitting circumstances based either partly or wholly on high probability (as per, e.g., Spohn 2002). Having identified the important properties of fitting circumstances, we want to leave open how exactly the further details are best specified. That said, throughout the paper we sometimes assume the dispositional account of *cp* laws, both for its *prima*

facie plausibility and for ease of exposition.

3.2 The role of *ceteris paribus* laws in supporting counterfactuals

Let us now show how *cp* laws, broadly interpreted as above, can underpin our proposal for understanding counterfactuals, or at least those for which the relevant domain is a special science. (Recall that our strategy is to convince the reader of our interpretation of these more obviously ‘scientific’ counterfactuals, such that the extension to more homely or ordinary counterfactuals is a natural move.) The good news is that many *cp* laws state deterministic relationships that may play a role in securing the truth of would-counterfactual claims. Moreover, the fact that *cp* laws restrict the circumstances to those supporting the regularity in question seems to fit well with our account of counterfactuals.

But now the challenge: one cannot simply take for granted that the antecedent of an ordinary counterfactual implies the very same circumstance restrictions that are implied by the pertinent *cp* law. Indeed, it is thought that the literal *inexactness* of *cp* law statements (the fact that, absent the implicit *ceteris paribus* clause, they are false) comes back to bite when it comes to counterfactuals.

After all, ‘shield volcanoes erupt effusively’ is only true because it does not apply to any old shield volcano; it applies to shield volcanos in fitting circumstances. But therein lies a worry: the antecedents of typical counterfactuals do not seem to be so carefully specified as to pick out the relevant fitting circumstances. Consider 5:

(5) Were Mt Pinatubo a shield volcano, it would have erupted effusively.¹²

The relevant *cp* law – ‘shield volcanoes erupt effusively’ – does not apparently make this counterfactual true, either because the antecedent is, in Hájek’s terms, *unspecific*, in encompassing all types of shield volcanos, or the connection between shield volcanos and effusive eruptions is

¹² Mt Pinatubo in the Philippines is a *stratovolcano*, as opposed to a *shield* volcano. It erupted explosively in 1991.

really *indeterministic*. In short, it appears as if the antecedent of 5 does not restrict to cases where shield volcanos erupt effusively.

We get around this problem in that we understand counterfactuals similarly to the statements of *cp* laws, as having an *implicit* ‘fitting circumstances’ clause in the antecedent, whose meaning varies with the salient domain of inquiry. Here we part ways with some proponents of the lawfulness of *cp* laws who do not rule out the line of reasoning expressed above regarding 5. Lange (2002), Woodward (2003) and Hitchcock and Woodward (2003), for instance, all maintain a close connection between *cp* laws and the truth of counterfactuals, but they claim that a counterfactual invoking a *cp* law is true *only if* the antecedent refers to the specific circumstances implicitly required by the *cp* law. It is not clear whether the specific circumstances should be taken to hold if this is not stated explicitly. Lipton’s (1999, 157) position is clearer; when it comes to counterfactuals, or indeed any manifestation of the disposition expressed by a *cp* law, it is not assured that there will be no interfering factors, which complicates the usual relationship between lawful regularities and the way things are or would be:

‘‘All Fs are Gs, *cp*’ may be a law yet not entail that if something had been an F it would have been a G, nor will observed Fs that are G always provide reason to believe that the next F will be a G as well, since we may have no reason to believe that all things will be equal, the next time.’

According to this view, *cp* laws do not, after all, secure the truth of would-counterfactuals. Statements like 5 are false, since, on closer analysis, Mt Pinatubo *might* be a pyroclastic shield volcano.

Despite the presumption in favour of the *inexact* reading of would-counterfactuals, we maintain that this asymmetry – between carefully circumscribed *cp* law statements and uncircumscribed counterfactuals which invoke them – should be resisted. The asymmetry is puzzling in light of what seems to be common ground for *all* defenders of *cp* laws – that *cp* laws are regularities pertaining to particular domains of inquiry involving particular sorts of entities under fitting circumstances. Plausibly then, if *cp* laws are supposed to play a role in determining whether at

least some counterfactuals are true, then the counterfactuals in question must *already* be understood as concerning the relevant domain of inquiry. In that case, one would expect the antecedent of the counterfactual to be understood as involving the parallel qualifications tied to that domain.

We regard it an attractive feature of our proposal that it allows counterfactuals to serve as tools for making claims whose truth values are tied to the many and varied regularities in the actual world. Some comments on this point are in order, although we save a more detailed analysis for future work. On our proposal, counterfactuals are true in virtue of facts about classes of circumstances that include more or fewer kinds of interferences, depending on the domain of inquiry. To be sure, sometimes the class of circumstances may include all the many and varied interferences that are physically compatible with the antecedent as stated (the trivial case where the domain is the most fundamental one; recall footnote 4). For instance, 5 may be read in this way if the conversation is one where a committee is pondering the extent to which a public official is responsible for the devastating effects of Mt Pinatubo's eruption, there having been reasonable uncertainty about the type of volcano it was. On this reading, 5 is false: Mt Pinatubo may have been an unusual shield volcano that did not erupt effusively.

We suggest, however, that it is less common than one might suppose that counterfactuals need be interpreted such that all possible interferences are under consideration. Often, we maintain, the counterfactuals pertinent to a conversation simply express interesting, constrained regularities in the actual world. Take 5 above. On an alternative and very natural reading, the counterfactual conveys what the eruption would have looked like, were Mt Pinatubo an ordinary shield volcano, that is, under the circumstances fitting to ordinary volcano science. In this case, 5 is true: circumstances in which Mt Pinatubo is an unusual shield volcano that erupts explosively count as interferences and are thus excluded.

In other cases, counterfactuals are used to indirectly explain actual events by stating what would have happened if things had been different. But any such (complete or partial) explanation of an actual happening relies on the lawful regularities that hold across a pertinent class of cases. For instance, consider 5*:

(5*) Were Mt Pinatubo a shield volcano, it would have erupted effusively and had a less dramatic aftermath for the inhabitants of the area.

This counterfactual plausibly conveys information about Mt Pinatubo, specifically the impact of its eruption on the local inhabitants. The counterfactual provides a partial explanation of the volcano's impact by suggesting conditions under which the impact would have been different. We are effectively told that shield volcanoes under fitting conditions erupt less violently, and from this claim regarding basic volcano science are presumably supposed to infer that Mt Pinatubo is not a shield volcano, and that is why its eruption had the dramatic impact it did.

Spelling out a precise taxonomy of the scientific claims that may be expressed by counterfactuals is beyond the scope of this paper. The point is simply that our appeal to domains of inquiry and the inclusion of an implicit 'fitting circumstances' proviso in the antecedent of counterfactuals is well-motivated, based on the relationship between counterfactuals and the diverse and domain-specific nature of scientific claims.

4. Comparison with *near-miss* Theories

One might be concerned that there are simpler proposals that have the same advantages. In particular, so-called *near-miss* proposals apparently fit this bill.¹³ Common to these proposals is the idea that, though there are cases in which the antecedent of a given counterfactual is true while its consequent is false, the counterfactual may nonetheless be true if these cases are objectively *too remote*.

One common way to understand what it is to be remote is in terms of probability: an objectively remote event is one that has sufficiently low probability. Let us first consider a simple near-miss view based on objective conditional probabilities according to which 'if P were true then Q would

¹³ Some examples of near-miss proposals are those defended in Bennett (2003, 246-259), Williams (2008), and Leitgeb (2012).

be true' ($P \square \rightarrow Q$) is true just in case the probability of Q given P is sufficiently high. This is a simplified version of the account Leitgeb (2012) develops. Bennett (2003, 246-259) and Williams (2008) also adopt near-miss approaches, though the central idea is implemented in different ways. We focus on the simple proposal and connect it up with some of the other near-miss proposals below.

The simple near-miss proposal yields a response to counterfactual scepticism. If the probability that the glass shatters given that you drop it is high enough, then 1a will come out true. The troubling possibilities to which the sceptic appeals are typically improbable, and thus, need not threaten the truth of would-counterfactuals.

Any gain in simplicity, however, is undone by two challenges for near-miss proposals. In what follows, we discuss these challenges in turn, one concerning lotteries and agglomeration (4.1), and the other concerning counterfactual modus ponens (4.2).¹⁴

4.1 Lotteries and Agglomeration

Consider a 1000-ticket lottery with one prize that is never played. 6 rings false.

(6) If I bought a ticket, I would have lost.

I *might not* have lost! However, the probability that I would have lost given that I buy a given ticket is high.¹⁵ So the simple near-miss proposal delivers the wrong result that 6 is true in the situation described.

¹⁴ These inferences may be invalidated for other reasons (see Lycan (2001), Briggs (2012), and Santorio (2019)). But we need not invalidate these inference patterns *in virtue of our response to counterfactual scepticism*.

¹⁵ If you think 1:1000 odds ratio is too high, such that the probability of a ticket losing is not sufficiently high, make it a million-ticket lottery.

A related challenge for near-miss proposals concerns an argument form often called *agglomeration* (see Hawthorne (2005), Hájek (ms), and K. Lewis (2016)):¹⁶

$P \square \rightarrow Q_1, P \square \rightarrow Q_2$, therefore $P \square \rightarrow (Q_1 \ \& \ Q_2)$

Consider the following series:

(7i) If the 1000-ticket lottery were played, ticket 1 would have lost.

(7ii) If the 1000-ticket lottery were played, ticket 2 would have lost.

...

(7m) If the 1000-ticket lottery were played, ticket 1000 would have lost.

On the simple near-miss proposal, (7i) through (7m) are all true.¹⁷ And yet (7*) must be false, since the probability of the consequent given the antecedent is zero.

(7*) If the 1000-ticket lottery were played, ticket 1 would have lost *and* ticket 2 would have lost... *and* ticket 1000 would have lost.

This is a violation of agglomeration.

Our proposal delivers the correct results in these cases. 6 and 7i through 7m are false; any particular lottery ticket winning, no matter how improbable, occurs under fitting circumstance for the salient domain of inquiry (plausibly, the *science of lotteries*). So there is no violation of agglomeration and we deliver the right verdict about 6. In general, on our account, the ‘agglomerated’ counterfactual either implies the same domain of inquiry and so has the same truth value as the conjunction of the counterfactuals in the sequence, or the ‘agglomerated’ counterfactual implies a different domain of inquiry from the counterfactuals in the sequence. Either way, there is no violation of agglomeration.

¹⁶ We are taking ‘argument forms’ to apply at the level of *propositions* rather than sentences; in the characterisation of the argument form, ‘P’ and ‘Q’ each stand for some proposition.

¹⁷ Again, if you think the 1:1000 odds ratio is too high, make it a million-ticket lottery.

4.2 Counterfactual Modus Ponens

Let's now consider counterfactual modus ponens ($P, P \square \rightarrow Q$, therefore Q). Suppose 1a is true and you drop the glass and it does not shatter in virtue of an improbable occurrence. This constellation of circumstances is possible given the simple near-miss proposal, but this is a case in which P and $P \square \rightarrow Q$ are true, yet Q is false, invalidating counterfactual modus ponens.

Our proposal might also seem to incur this same cost. Not so! Suppose someone utters 1a, pointing to a glass I am holding. Suppose further that I drop the glass and there is some aberrant physical arrangement that stops the glass from shattering. The presence of this arrangement either counts as a fitting circumstance or it does not. If it does not, the counterfactual involving fitting circumstances may be true but the antecedent, interpreted as we propose, is not actually satisfied. If the arrangement does count as a fitting circumstance, then the counterfactual is false. In either case, the validity of counterfactual modus ponens is not threatened.¹⁸

There is a general issue about what happens to the truth values of counterfactuals when strange, unlikely, or remarkable events occur at the actual world. Consider:

(8) Were you to drink heavily tonight, you would be hungover tomorrow morning.

Many are inclined to judge that this claim is true. It expresses something important about human biology and alcohol.

¹⁸ A defender of a near-miss proposal could avoid invalidating counterfactual modus ponens and agglomeration by helping themselves to the resources we offer. For instance, they could claim that the threshold for 'high enough probability' shifts with the implied domain of inquiry or that the implied domain of inquiry makes a difference to the correct interpretation of the antecedent. But once one helps oneself to the resources we offer, what motivation is there to *also* go in for a near-miss-based response to counterfactual scepticism?

But suppose that at 11:59pm tonight the sun expands and engulfs the earth, there is no morning, let alone a hangover. Some are, like us, inclined to suggest that there is nonetheless a true reading of 8; there is still an important lawful connection between drinking and hangovers that one good reading of 8 expresses. Many who, like Briggs (2012) and Lycan (2001), offer a true reading of 8 are prepared to give up on counterfactual modus ponens. Accordingly, it might seem as if we must either give up counterfactual modus ponens or give up the intuition that there is an acceptable true reading of 8 in the world in which the sun expands.

Our proposal allows this intuition to be vindicated while avoiding invalidating counterfactual modus ponens. In line with our proposal, we propose interpreting 8 as 8s.

(8s) Were you to drink heavily tonight *under fitting circumstances for domain d*, you would be hungover tomorrow morning.

If what we are interested in expressing when uttering 8 is the lawful regularity between drinking and hangovers, then we will not consider the sun-expansion scenario to be a fitting circumstance from the point of view of the relevant domain of inquiry, *d*. So 8 may be true, without violating counterfactual modus ponens, even if the agent in question drinks heavily and the earth is engulfed before midnight, since the latter event means that the actual world turns out to *not* be one in which the antecedent, interpreted in line with the relevant domain of inquiry, holds.

We have considered how a simple near-miss proposal faces the challenges of agglomeration and counterfactual modus ponens. Given these prima-facie challenges, defenders of near-miss proposals can, of course, complicate their proposals to either validate these inference patterns within their framework or explain why they should be rejected in any case.¹⁹ However, recall that the purpose of this section is not to show that near-miss proposals should be rejected in the face of these challenges. Rather our goal was to show that, insofar as near-miss accounts are simpler alternatives to our proposal, they do *not* have the same advantages.

¹⁹ Recall the suggestion in footnote 18. See too Leitgeb (2012, 109-117) on agglomeration and Williams (2008, 394-396) on counterfactual modus ponens.

5. Shifts

Complexities in ordinary counterfactual talk have historically been a focal point for adjudicating between competing approaches to counterfactuals. To illustrate, the unaugmented strict conditional analysis does not have the requisite flexibility to capture all the distinctions between counterfactuals that exist in ordinary language.

Recall that the strict conditional analysis says that $P \Box \rightarrow Q$ is true just in case $\Box(P \supset Q)$. This analysis runs afoul of ‘Sobel sequences’ like the following:

(9a) Were Alice to come to the party, it would be fun.

(9b) Were Alice to come to the party and there was a major house fire, it would not be fun.

9a rings true but so does 9b. The unaugmented strict conditional analysis cannot allow for this pattern of truth values. The reason is that if a statement of the form, $\Box(P \supset Q)$ is true, then conjoining the antecedent with any further proposition R cannot make the resulting claim false (or equivalently, it cannot make $\Box((P \& R) \supset \sim Q)$ true for any R). Yet this is precisely what seems to occur in moving from 9a to 9b; the first rings true, but so does the second.

The ability to correctly handle sequences like this is a widely recognised desideratum for theories of counterfactuals. Our proposal can deliver the correct verdicts about Sobel sequences. We will get to the details below. For now, let us table another kind of sequence that has proved challenging: ‘Heim sequences’.²⁰ Consider the following rearrangement of 9a and 9b:

(9b) If Alice were at the party and there was a major house fire, it would not be fun.

(9a) If Alice were at the party, it would be fun.

Where 9a is entertained immediately after entertaining 9b, 9a may reasonably be resisted. After

²⁰ Named for Irene Heim.

all, when entertaining 9a immediately after 9b, it is hard not to focus on cases in which Alice is at the party and there is a house fire.²¹ But this seems to be in direct tension with the previous judgment according to which 9a and 9b are both true.

The puzzle introduced in section 1 has a similar flavour. While it seems as if statements like 1a and 1b cannot both be true, it is also apparently a feature of ordinary reasoning that one can correctly assert a would-counterfactual, and yet, at least in some cases, subsequently concede the truth of the corresponding might-not-counterfactual. How can a proposal for understanding counterfactuals satisfy these apparently conflicting demands?

One response is to side-line these puzzles to the domain of conversational pragmatics: what does and does not sound right in these contexts need not track the truth or falsity of the propositions expressed.²² Others, us included, seek a semantic treatment of these puzzles.²³ For this latter strategy to work, what is required is the further versatility afforded by an extra *contextual* parameter governing the semantics of counterfactuals. Those who have pursued this line tend to suggest that the extra parameter tracks ‘conversational context’ (von Stechow 2001, Gillies 2007, Ichikawa 2011, Lewis 2016).

According to the recent and elegant proposal of K. Lewis, for instance, ‘counterfactuals are not just evaluated relative to the most similar antecedent worlds, but relative also to the worlds relevant given the context’ (2016, 2). That is, irrelevant possibilities are either excluded from the set of

²¹ We do not mean to claim that all Heim sequences will be infelicitous.

²² Moss (2012, 2013), for instance, offers a sophisticated version of this line of response.

²³ Emery (2017) denies such a sharp distinction; she claims that whether a response to counterfactual scepticism is best construed as a semantic or as a pragmatic proposal depends on how one prioritises ‘fairly abstract theoretical virtues that one’s semantics and pragmatics can exhibit’ (409). We are sympathetic to this line, but consider the semantic route to have at least expository advantages; for one thing, desiderata (such as Emery’s own ‘objectivity constraint’, which is akin to K. Lewis’s notion of ‘objectivity’, discussed below) are typically more readily articulated for semantic proposals.

worlds under consideration or are characterized as *distant* possibilities. This yields a response to counterfactual scepticism, because the possibilities to which the sceptic appeals may be irrelevant to the conversational purposes in play. There can also be shifts in conversational context, yielding shifts in the truth value of a given counterfactual. K. Lewis tells a plausible story as to how it works for the sorts of examples mentioned above and in section 1, and her theory allows for the truth values of counterfactuals to shift accordingly.²⁴

Our proposal also has a certain amount of flexibility, given that it too has an extra contextual parameter: the contextually salient domain of inquiry, which governs which lawful regularities are at issue. The domain of inquiry can shift and so change the interpretation of a counterfactual, as per the transition from 1a to 1b. But there is only limited leeway for such shifts, whatever the complete story about the right domain for interpreting a counterfactual utterance.²⁵ For one thing, the domain is to some extent dependent on the wording of the counterfactual. As noted above, 1a is more naturally interpreted with respect to the domain of macro-object dynamics than statistical mechanics. It requires weighty conversational cues to shift the domain against the wording (as per the shift from 1a to 1b given a persistent interlocutor). Moreover, the conversational cues may themselves be quite complex and resistant to shifts. For instance, if the conversation suggests that a counterfactual *explain* an actual happening, the candidate domains may be quite limited, even before other considerations are taken into account.

Sobel and Heim sequences draw attention to how the wording or surface form of the antecedent affects a counterfactual's interpretation. On our proposal, the transition from 9a to 9b may involve a shift in the domain at play due to the more detailed antecedent – a shift from how Alice behaves at parties to, say, how groups including Alice respond to a broad range of threats (since the surface

²⁴ As noted earlier in footnote 5, K. Lewis defends contextualism about counterfactuals in the face of contrary linguistic data that suggests a semantic treatment of the aforementioned puzzling sequences is not the right way to go. We will not rehearse Lewis's arguments here, but we of course welcome her general defence of contextualism.

²⁵ Recall that we do not pretend to give the full story here. We merely illustrate how such a story can resolve well-known puzzles about counterfactuals.

form of the antecedent in 9b suggests that fewer events should be treated as interferences by explicitly mentioning one such event). Different things will count as fitting circumstances on the corresponding interpretations of 9a and 9b. This captures how 9a and 9b can be true together. Nonetheless, if 9b is uttered first, the domain of inquiry arguably does not shift upon the introduction of 9a into the conversation. Once the possibility of a house fire has been introduced into the conversation it is harder to treat events like house fires as interfering factors. What is more, this appears to be true even when we drop the explicit mention of the house fire from the surface form of the antecedent, as per the second sentence in the Heim sequence. Just as epistemic standards (e.g. requirements for knowledge) are often easy to raise but harder to lower, it seems that it is easier to make some possibility salient than to render that same possibility non-salient (such that it can be treated as an interference) once it has been explicitly mentioned. We suggest that this is why our account, as per other forms of contextualism, can explain why shifts are often more natural in Sobel sequences than in Heim sequences.

The most important feature that sets our proposal apart from other context-sensitive accounts of counterfactuals is that we foreground the *objectivity* of the facts that make true counterfactuals true. If K. Lewis is right, facts about counterfactual dependence seem to rock and sway with what groups of humans are interested in talking about. One might worry that the set of facts about what counterfactually depends on what is not so changeable and tying counterfactual dependence to conversational purposes compromises the status of counterfactuals as genuinely worldly statements about the way things are. Note that this is not merely a worry about radical subjectivity. K. Lewis convincingly staves off the objection that, on her proposal, a counterfactual is true just in case the speaker holds that it is true for her own idiosyncratic conversational purposes. According to K. Lewis, conversational context has a certain objectivity, such that speakers may themselves be misguided about the conversational context in which they find themselves.²⁶ But

²⁶ Indeed, K. Lewis goes further in suggesting that conversational contexts may describe social institutions, as it were, rather than personal discussions. She goes so far as to suggest that different scientific disciplines may count as conversational contexts. We obviously have a lot of sympathy for this way of conceiving conversational contexts. But our proposal draws a stronger and more steadfast connection between counterfactuals and science. We do not see science as just another

her proposal still allows that the local circumstances of human thought and talk, however objectively these circumstances may be described, affect the truth values of counterfactual propositions.

On the other hand, we propose that the facts about counterfactual dependence are fixed by the facts about regularities in the world.²⁷ There are just more such facts out there than one might have expected, and it is a subtle business as to which of these facts is pertinent to a given counterfactual utterance or thought. Once we have identified the appropriate domain of inquiry and corresponding interpretation of a counterfactual statement, it is an *objective* matter whether that counterfactual, thus interpreted, is true.

6. Concluding Remarks

Hájek begins his defence of counterfactual scepticism by distinguishing two camps of people interested in counterfactuals: those who think that counterfactual scepticism cannot be the right diagnosis of our ordinary discourse (predominantly philosophers of language) and those who think counterfactual scepticism must be right given the important sense in which our world is chancy (predominantly philosophers of science). With those in the ‘philosophy of science’ camp, Hájek contends that taking the role of counterfactual claims in science seriously leads one to counterfactual scepticism. We have argued that this is not the case. Science, understood as a diverse, textured, and multi-leveled enterprise, leaves an important place for true, substantive, and objective, non-probabilistic ‘would’ counterfactuals. We have also given reasons to think that our proposal fits well with the considerations of interest to the ‘philosophy of language’ camp; our account allows us to rise to the challenge of the complexity of ordinary counterfactual thought and

conversational context. Science is systematic inquiry about Nature’s many domain-specific regularities, and, we suggest, it is these regularities that govern the truth of counterfactuals.

²⁷ We allow that the collective aims or values underpinning the institutions of science may influence what count as ‘scientific laws’. But we need not take a definitive stance on this issue. Either way, the truth of counterfactuals does not depend on the idiosyncratic details of any localised conversational context.

talk without sacrificing objectivity. Our proposal suggests a method of harmonising constraints stemming from philosophy of science with those stemming from philosophy of language, allowing us to meet Hájek’s sceptical challenge without compromise. This is reason enough to take our proposal seriously.²⁸

References

- Bennett, Jonathan (2003) *A Philosophical Guide to Conditionals*. Vol. 114: Oxford University Press.
- Briggs, Rachael (2012) “Interventionist counterfactuals”, *Philosophical Studies* 160 (1):139-166.
- Cartwright, Nancy (1989) *Nature's Capacities and their Measurement*, Cambridge: Cambridge University Press.
- Emery, Nina (2017) “The Metaphysical Consequences of Counterfactual Skepticism”, *Philosophy and Phenomenological Research* 94(2): 399–432.
- von Fintel, Kai (2001) “Counterfactuals in a Dynamic Context”, in Ken Hale: A Life in Language, Michael Kenstowicz (ed.) Cambridge: MIT Press.
- Fodor, Jerry (1987) *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, Cambridge, MA: MIT Press.
- Gillies, Thony (2007) “Counterfactual Scorekeeping”, *Linguistics and Philosophy* 30: 329-360.
- Goodman, Nelson (1955) *Fact, Fiction, and Forecast*. Vol. 31: Harvard University Press.
- Hausman, Daniel M. (1992) *The Inexact and Separate Science of Economics*. Cambridge: Cambridge University Press.
- Hawthorne, John (2005) “Chance and Counterfactuals”, *Philosophy and Phenomenological Research* 70(2): 396-405.
- Hájek, Alan. ms. “Most Counterfactuals are False”.

²⁸ Many thanks to the audiences of the Fourth Taiwan Philosophical Logic Colloquium (2018), the Causal Models and Counterfactuals Workshop at ANU (2018), the Melbourne Logic Seminar (2017) and the AAP in Adelaide (2017) for helpful discussions on the content of this paper. In addition to the editors of this special issue and the two anonymous reviewers, we are particularly grateful to Renée Jorgensen Bolinger, Alan Hájek and Daniel Nolan for extensive comments on earlier drafts of the paper. Research for this paper was supported by Katie Steele’s ANU Futures Grant and Alexander Sandgren’s Vetenskapsrådet Research Project Grant (2019-02786).

- Hitchcock, Christopher, and Woodward, James (2003) “Explanatory Generalizations, Part II: Plumbing Explanatory Depth”, *Noûs*, 37(2): 181–199.
- Ichikawa, Jonathan (2011) “Quantifiers, Knowledge, and Counterfactuals”, *Philosophy and Phenomenological Research* 82(2): 287–313.
- Kincaid, Harold (2004) “Are There Laws in the Social Sciences?: Yes”, in *Contemporary Debates in the Philosophy of Science*, C. Hitchcock (ed.), Oxford: Blackwell, pp. 168–187.
- Kratzer, Angelika (2012) *Modals and Conditionals: New and Revised Perspectives*, vol. 36. Oxford University Press.
- Kutach, Douglas (2005) “Counterfactuals in Science”, *Encyclopedia of Philosophy*. Encyclopedia.com. 29 Nov. 2017 <<http://www.encyclopedia.com>>.
- Lange, Marc (2000) *Natural Laws in Scientific Practice*, Oxford: Oxford University Press.
- Lange, Marc (2002) “Who's Afraid of Ceteris Paribus Laws? Or: How I Learned to Stop Worrying and Love Them”, in *Ceteris Paribus laws*, J. Earman, et al. (eds.), *Erkenntnis*, 52 (Special Issue): 407–423.
- Leitgeb, Hannes (2012) “A probabilistic semantics for counterfactuals”, *The Review of Symbolic Logic* 5(1): 26–121.
- Lewis, David K. (1973) *Counterfactuals*: Blackwell.
- Lewis, David K. (1979) “Counterfactual Dependence and Time's Arrow”, *Noûs* 13: 455–76. Reprinted with postscript in Lewis, *Philosophical Papers II* (Oxford University Press, 1986) 32–51. Also reprinted in Jackson (ed) *Conditionals* (Oxford University Press, 1991) 46–76.
- Lewis, Karen S. (2016) “Elusive Counterfactuals”, *Noûs* 50(2): 286–313.
- Lipton, Peter (1999) “All Else Being Equal”, *Philosophy* 74: 155–168.
- List, Christian and Pivato, Marcus (2015) “Emergent Chance”, *The Philosophical Review* 124(1): 119–152.
- Lycan, William G (2001) *Real Conditionals*, Oxford, UK: Oxford University Press.
- Mill, John S., (1836/2008) “On the Definition and Method of Political Economy”, in *The Philosophy of Economics. An Anthology*, D. Hausman (ed.), New York: Cambridge University Press, Third Edition, pp. 41–58.
- Moss, Sarah (2012) “On the Pragmatics of Counterfactuals”, *Noûs* 46 (3): 561–586.

- Moss, Sarah (2013) “Subjunctive Credences and Semantic Humility”, *Philosophy and Phenomenological Research* 87(2): 251–278.
- Reutlinger, Alexander, Schurz, Gerhard and Hüttemann, Andreas (2017) “*Ceteris Paribus* Laws”, *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2017/entries/ceteris-paribus/>>.
- Roberts, John T. (2004) “There are No Laws in the Social Sciences”, in *Contemporary Debates in the Philosophy of Science*, C. Hitchcock (ed.), Oxford: Blackwell, pp. 168–185.
- Rosenberg, Alex and McShea, Daniel W. (2008) *Philosophy of Biology: A Contemporary Introduction*. New York: Routledge.
- Santorio, Paolo (2019) “Interventions in Premise Semantics”, *Philosophers Imprint*, 19 (1): 1-27.
- Schulz, Moritz (2010) “Counterfactuals in a Chancy World”, ms. Department of Philosophy, Oxford University.
- Spohn, Wolfgang (2002) “Laws, *Ceteris Paribus* conditions, and the Dynamics of Belief”, in *Ceteris Paribus laws*, J. Earman, *et al.* (eds.), *Erkenntnis*, 52 (Special Issue): 373–394.
- Stefánsson, H. Orri (2018) “Counterfactual Skepticism and Multidimensional Semantics”, *Erkenntnis* 83(5): 875–898.
- Stalnaker, Robert C. (1968) “A Theory of Conditionals”, in *Studies in Logical Theory, American Philosophical Quarterly* (Monograph Series, 2), Oxford: Blackwell, pp. 98–112.
- Strevens, Michael (2014) ‘High Level Exceptions Explained’, *Erkenntnis* 79:1819–1832.
- Williams, Robert (2008) “Chances, Counterfactuals, and Similarity”, *Philosophy and Phenomenological Research* 77 (2): 385–420.
- Woodward, James (2003) *Making Things Happen*, Oxford: Oxford University Press.