Counterfactuals of Ontological Dependence

Abstract

A great deal has been written about 'would' counterfactuals of causal dependence. Comparatively little has been said regarding 'would' counterfactuals of ontological dependence. The standard Lewis-Stalnaker semantics is inadequate for handling such counterfactuals. That's because some of these counterfactuals are counterpossibles, and the standard Lewis-Stalnaker semantics trivializes for counterpossibles. Fortunately, there is a straightforward extension of the Lewis-Stalnaker semantics available that handles counterpossibles: simply take Lewis's closeness relation that orders possible worlds and unleash it across impossible worlds. To apply the extended semantics, an account of the closeness relation for counterpossibles is needed. In this paper I offer a strategy for evaluating 'would' counterfactuals of ontological dependence that understands closeness between worlds in terms of the metaphysical concept of grounding.

Keywords Counterfactuals, Ontological Dependence, Grounding, Metaphysical Laws, Counterpossibles

1. Introduction

When the existence of one entity, $e*$, depends on the existence of another, e, we say that $e*$ is ontologically dependent upon e. Ontological dependence is distinct from causal dependence. Causal dependence is a relation between events across time; ontological dependence is a relation between entities that exist at a time, or between entities that lack a spatiotemporal location altogether. For both kinds of dependence there are true 'would' counterfactuals. For instance, if Suzy throws a rock at a window, causing it to break, the counterfactual 'if Suzy had not thrown the rock, the window would not have broken' appears to be true. But by the same token, if the existence of the singleton set $\{2\}$ wholly depends upon the existence of its urelement, the counterfactual 'if 2 had not existed, $\{2\}$ would not have existed' appears to be true.

The standard Lewis-Stalnaker semantics is inadequate for handling all 'would' counterfactuals of ontological dependence. That's because some such counterfactuals concern entities that exist necessarily, as is the case for the number 2 and its singleton, and so are counterpossibles. The trouble is that the standard Lewis-Stalnaker semantics trivializes for counterpossibles. Fortunately, there is a straightforward extension of the Lewis-Stalnaker semantics available that handles counterpossibles: simply take Lewis's closeness relation that orders possible worlds and unleash it across impossible worlds.

Very little has been said about the similarity relation that underpins this extension of the Lewis-Stalnaker semantics. In this paper, I address this lacuna by offering an account of the similarity relation for counterfactuals of ontological dependence. The account is tailored toward counterfactuals with impossible antecedents, but it applies to counterfactuals with possible antecedents as well. Note that my goal is just to provide an account of particular judgements of specific ontological dependencies. I am not trying to provide anything like a counterfactual analysis of ontological dependence in general, or of the metaphysical laws that govern it (if there are any such laws, more on this later). Indeed, any such analysis along these lines using the account developed here would most likely be circular. The notion of a metaphysical law (or as I call it, a grounding law) and the more general notion of grounding are used to determine the similarity between worlds. Since there is a deep metaphysical connection between grounding and ontological dependence, and between ontological dependence and metaphysical laws more generally, a counterfactual analysis of such notions using the picture presented here would presuppose the very phenomena being analysed.

I begin by providing a brief recap of the Lewis-Stalnaker semantics for counterfactuals and its extension to counterpossibles (§2), before saying a bit about why it is important to provide an account of the similarity relation that appears in this analysis, and a bit about the methodology (§3). I go on to offer an account of the similarity relation for the 'would' counterfactuals at issue that understands closeness between worlds in terms of grounding (§4). I then demonstrate the account on a toy model (§5). The model represents a range of ontological structures for which 'would' counterfactuals of ontological dependence arise and so provides evidence for the wider applicability of the account. I wrap up in §6.

2. Semantics

Call any counterfactual with the following broad form a 'would' counterfactual of ontological dependence: if e had not existed, $e*$ would not have existed (where e and $e*$ either exist at the same time, or lack a spatiotemporal location entirely). Now, consider the following 'would' counterfactual that accompanies the ontological dependence of $\{2\}$ on its urelement:

> (CP_1) If the number 2 had not existed, the singleton set $\{2\}$ would not have existed.

If numbers and sets exist then (CP_1) seems to be true. On the reasonable assumption that mathematical objects exist necessarily it follows that (CP_1) is a counterpossible. The standard Lewis-Stalnaker semantics for counterfactuals cannot handle counterpossibles. Lewis's (1973: 424–425) semantics for counterfactuals may be stated as follows:

> **Analysis 1** $A \mathrel{\Box\!\!\rightarrow} B$ is true at a world ω iff some possible world in which both A and B are true is closer to ω than any possible world in which A is true and B is false, if there are any possible worlds in which A is true.

Closeness is a measure of comparative similarity and is subject to (at least) two constraints. First: *weakness*. A closeness ordering of worlds permits possible worlds to tie for closeness, but any two possible worlds are comparable. Second: *centering*. The actual world is the closest possible world to itself, resembling itself more than any other possible world.

Analysis 1 renders all counterpossibles vacuously true. For instance, consider (CP_1). Numbers are metaphysically necessary entities. So there is no world in which the number 2 does not exist. So the statement on the right-hand side of the biconditional in **Analysis 1** is trivially satisfied. This might be fine except the following counterfactual is also true, according to **Analysis 1**:

(CP_{1*}) If the number 2 had not existed, my shoes would not have existed.

(CP_{1*}) is outrageous; the number 2 has nothing to do with my shoes. To avoid vindicating *both* (CP_1) and (CP_{1*}) Lewis (1973: 434) notes that the semantics may be made to range over both possible and *impossible* worlds (though he resists the urge to do so). Extending the semantics in this fashion yields the following (Beall and van Fraassen (2003); Laan (2004); Mares (1997); Nolan (1997) and Restall (1997)):

> **Analysis 2** $A \mathrel{\Box\!\!\rightarrow} B$ is true at a world ω iff some possible or impossible world in which both A and B are true is closer to ω than any possible or impossible world in which A is true and and B is false, if there are any possible or impossible worlds in which A is true.

As before, closeness is subject to *weakness* and *centering*.

Analysis 2 helps because a counterfactual with an impossible antecedent may have a true antecedent in some impossible world. Accordingly, such a counterfactual may not be vacuously true when evaluated against both possibility and impossibility. Indeed, such a counterfactual may not be true at all. Using **Analysis 2**, then, there's scope to make (CP_1) true without also making (CP_{1*}) true. Of course, while Lewis suggests this extension of his own semantics, his considered view is that counterfactuals with necessarily true antecedents really are vacuously true, and thus his original semantics gets the right results. Williamson (2013) has recently provided a defense of this line of thought. I will not engage in this debate here, except to note that there are now a number of compelling responses to Williamson's arguments available (see, for instance, Berto et al. (forthcoming)). If one sides with Williamson and Lewis and thinks that all such counterfactuals are vacuously true, then one need read no further. If one believes as I do that Williamson and Lewis are wrong, and that the weight of evidence is against them, then read on.

What are impossible worlds? Mares (1997), Priest (2002) and Restall (1997) offer the same basic answer to this question: whatever the correct metaphysical account of possible worlds might be, that account can be carried over to impossible worlds. According to Jago (2015), however, not every metaphysical account of possible worlds can be generalized in the manner imagined. But all that **Analysis 2** requires is the weaker claim that there is at least one metaphysical account of possible worlds that applies to impossible worlds as well. This weaker claim is plausible; Jago (2015) offers one such account, Ripley (2012) offers another.

For now, I will sidestep this issue and simply assume a linguistic ersatz account of worlds, according to which a world is a set of sentences in an appropriate world-building language (such as Lagadonian sentences). According to such an account, a possible world can be defined as follows:

> ω **is a possible world** $=_{df}$ ω is a set S of sentences $P_1 ... P_n$ in a world-building language that is (i) maximal and (ii) consistent.

Note that by 'consistent' I mean 'negation consistent', where a set of propositions is negation consistent when for any proposition A it is not the case that both A and $\neg A$ are members of that set. Note also that by 'maximal' I mean to invoke bivalence. Thus, a set of propositions is maximal when for any proposition A and its negation $\neg A$, at least one of those propositions is in the set.

Following Bjerring (2014: 333–334), the truth and falsity of a sentence with respect to a world can be defined as follows:

> [**Truth**] A sentence P is *true* in a world w iff $P \in w$.

> [**Falsity**] A sentence P is *false* in a world w iff $P \notin w$.

Maximality is then defined as follows:

> [**Maximality**] A set Γ of sentences is *maximal* iff for all sentences P, either $P \in \Gamma$ or $\neg P \in \Gamma$ and not both.

Using the above definition of a possible world we can then define an impossible world as follows:

> ω **is an impossible world** $=_{df}$ ω is a set S of sentences $P_1 ... P_n$ in a world-building language that is either (i) non-maximal or (ii) inconsistent.

However, this definition of an impossible world is a bit too strong for what I have in mind; every impossible world is either inconsistent or non-maximal. When evaluating 'would' counterfactuals of ontological dependence it does not seem right to perform the evaluation against inconsistent or non-maximal impossibilities. For instance, suppose that the number 2 exists and that it does so of necessity. Imagining what would be the case were that number not to exist (as we must do when considering (CP_1)) does not require considering any inconsistent or incomplete worlds. Rather, we need only consider impossible worlds where numbers do not exist. Such worlds will differ from any possible world in that mathematical nominalism – as opposed to Platonism – is true, but should be perfectly consistent and maximal for all that.

So the definition of an impossible world must be expanded. Inspired by Nolan's (1997: 543–544) discussion of a similar issue, I propose the following refinement:

> ω **is an impossible world** $=_{df}$ ω is a set S of sentences $P_1 \ldots P_n$ in a world-building language such that either (i) S is inconsistent or (ii) S is non-maximal or (iii) there is at least one sentence $P \in S$ such that P is assigned a truth-value that no possible world assigns to P.

Note that what it means to say that there is at least one sentence $P \in S$ such that P is assigned a truth-value that no possible world assigns to P is that there is at least one sentence $P \in S$ such that there is no possible world that has P as a member or, equivalently, every possible world has $\neg P$ as a member.

The crucial clause is (iii). Suppose that '2 exists' is necessarily true. Given the above definition, there is a perfectly consistent and maximal world that is nevertheless an impossible world, simply because it holds that '2 exists' is false. Call a world that satisfies clauses (i) or (ii) a *nasty* impossible world and a world that satisfies clause (iii) but not clauses (i) or (ii) a *nice* impossible world. The counterfactuals in which I am interested require the consideration of nice, and not nasty, impossibilities. Accordingly, in what follows I will restrict **Analysis 2** to possible worlds and nice impossible worlds.

3. Nixon's Revenge

So we have a candidate semantics for counterpossibles. On its own, however, **Analysis 2** only takes us so far. **Analysis 2** needs to be fleshed out with some account of the similarity relation it invokes. As Lewis puts the point with respect to **Analysis 1**:

> While not devoid of testable content – it settles some questions of logic – it does little to predict the truth values of particular counterfactuals in particular contexts. The rest of the study of counterfactuals is not fully general ... [Analysis 1] must be fleshed out with an account of the appropriate similarity relation, and this will differ from context to context. (Lewis (1979: 465)

An account of the similarity relation is an account of the weights that are given to various respects of similarity between worlds for a particular counterfactual (or class of counterfactuals) within a particular context (or class of contexts). With respect to **Analysis 2**, what I aim to provide is an account of the respects of similarity that matter for evaluating counterfactuals of ontological dependence in ordinary contexts. Broadly speaking, there are two ways to produce such an account. First, one might decide, independently and for philosophical reasons, which respects of similarity matter for ontological dependence, and then apply that to **Analysis 2**. Second, one might reverse engineer an account of the similarity relation from what we know about the truth or falsity of 'would' counterfactuals of ontological dependence such that, when combined with **Analysis 2**, the account gets the truth-conditions for those counterfactuals right.

Lewis cautions against the first way of proceeding:

> The thing to do is not to start by deciding, once and for all, what we think about similarity of worlds, so that we can afterwards use these decisions to test [Analysis 1]. What that would test would be the combination of [Analysis 1] with a foolish denial of the shiftiness of similarity. Rather, we must use what we know about the truth and falsity of counterfactuals to see if we can find some sort of similarity relation–not necessarily the first one that springs to mind–that combines with [Analysis 1] to yield the proper truth conditions. It is this combination that can be tested against our knowledge of counterfactuals, not [Analysis 1] by itself. In looking for a combination that will stand up to the test, we must use what we know about counterfactuals to find out about the appropriate similarity relation–not the other way around. (Lewis 1979: 467)

I will heed Lewis's advice and take the second approach noted above: I will start with the truth of counterfactuals and work back to an account of the similarity relation. Now, it might be thought that it is straightforward to provide such an account. The similarity relation for 'would' counterfactuals of ontological dependence, one might argue, is just this: the degree of closeness between two worlds is the degree to which the two worlds agree about what exists. The more that the ontologies of the two worlds overlap, the closer the two worlds are. The closest worlds are thus the worlds that have the most similar ontology to the actual world.

However, this simple account fails. A more sophisticated approach is needed. To see this, consider, again, (CP_1). As noted, (CP_1) seems to be true: remove the number 2 from a world, and you must also remove its singleton set. Suppose, however, we evaluate (CP_1) by considering whether a world in which the number 2 does not exist and its singleton $\{2\}$ does not exist is closer than a world in which 2 does not exist and $\{2\}$ exists anyway. The trouble comes this way: the singleton $\{2\}$ does a lot of work. It is responsible for the existence of the following infinite chain of sets: $\{\{2\}\}$, $\{\{\{2\}\}\}$, $\{\{\{\{2\}\}\}\}$, $\{\{\{\{\{2\}\}\}\}\}$ ... The singleton $\{2\}$ is also partially responsible for the existence of any set that has $\{2\}$ as a member, of which there is an infinite number as well.

Now, consider the following two worlds: ω_1 and ω_2. ω_1 and ω_2 match actuality exactly except in the following respects. In both worlds the number 2 does not exist.

In ω_1 the singleton $\{2\}$ does not exist and nor does the infinite chain of sets $\{\{2\}\}$, $\{\{\{2\}\}\}$, $\{\{\{\{2\}\}\}\}$, $\{\{\{\{\{2\}\}\}\}\}$ ... In ω_2, however, the singleton $\{2\}$ exists and so do the sets for which it claims responsibility. According to the simple account, ω_2 is more similar to the actual world than ω_1. After all, ω_2 is missing just one object (the number 2), whereas ω_1 is missing a whole slew of objects; it has lost a large fragment of the set-theoretic hierarchy. Thus, according to the simple account, there is some world in which both 2 and $\{2\}$ fail to exist that is further away from actuality than some world in which 2 does not exist and $\{2\}$ exists anyway. Of course, to show that (CP_1) is false, it must be demonstrated that every world in which both 2 and $\{2\}$ fail to exist is further from actuality than some world in which 2 does not exist and $\{2\}$ exists anyway. But that is straightforward to show (I leave the details to the reader).

The case just outlined is a direct analogue of Fine's (1975) Nixon case, and raises the same problem. The original Nixon case showed that a simple account of the similarity relation for 'would' counterfactuals of causal dependence won't do because it tends to make those counterfactuals false. Similarly, the simple account of 'would' counterfactuals of ontological dependence won't do because counterfactuals that are intuitively true, turn out to be false on that account. Lewis (1979: 467–475) responds to Fine by producing a complicated recipe for determining the closeness of worlds. In the next section I will lay out a similar recipe for 'would' counterfactuals of ontological dependence by modifying the basic ideas underlying the Lewisian picture (see Kment (2014) for a similar approach, based on explanation).

To be clear, by laying out a recipe for evaluating 'would' counterfactuals, I am *not* claiming that this recipe should be used as a method for evaluating *all* counterfactuals whatsoever, or even all counterfactuals with necessarily false antecedents. The recipe is for a particular class of counterfactuals, those implicated in cases of ontological dependence. One might balk at this: surely a semantics for counterfactuals ought to be fully general, handling 'would' counterfactuals of ontological dependence as a special case!

In a certain sense, however, the semantics *is* fully general. The underlying semantics is the closeness semantics specified in **Analysis 2**. I am assuming for the sake of argument that this semantics applies to all counterfactuals, and to 'would' counterfactuals of ontological dependence as a special case. What I deny is that the notion of similarity captured by 'closeness' is the same notion of similarity for every counterfactual. Different counterfactuals require different dimensions of similarity for their evaluation. This is why Lewis does not extend his recipe to all counterfactuals. Rather, he applies the recipe to a certain class of counterfactuals and is largely silent about the way we should think about similarity for other types of counterfactuals. For instance, Lewis does not expect his recipe to work for backtracking counterfactuals (see Lewis (1979: 458)) or for counterfactuals in which the antecedent and the consequent are not propositions about distinct events (see his response to Kim's (1973) in Lewis (1987)).

4. Grounding and Similarity

Let us begin with Lewis's (1979) four-step recipe for determining the closeness of worlds in ordinary contexts:

> [i] It is of the first importance to avoid big, widespread, diverse miracles.
> [ii] It is of the second importance to maximize the spatiotemporal region throughout which perfect match of particular fact prevails.
> [iii] It is of the third importance to avoid even small, localized, simple miracles.
> [iv] It is of little or no importance to secure approximate similarity of particular fact, even in matters that concern us greatly. (Lewis 1979: 474)

Where, for Lewis (1979: 469), miracles are defined inter-worldly:

> A miracle at ω_1, relative to ω_0, is a violation at ω_1 of the laws of ω_0, which are at best the almost-laws of ω_1. The laws of ω_1 itself, if such there be, do not enter into it.

As it stands, Lewis's recipe is inadequate for evaluating 'would' counterfactuals of ontological dependence, for two reasons. First, Lewis's definition of a miracle is typically understood to involve physical laws of nature. The physical laws of nature, however, do not appear relevant to evaluating 'would' counterfactuals of ontological dependence. Those counterfactuals seem to radically outstrip the physical laws, in this sense: change the laws as you like and it would still be the case that if 2 did not exist, then $\{2\}$ would not exist. Second, Lewis's recipe is specifically aimed at matching for spatiotemporal regions. When evaluating counterpossibles regarding necessarily existing entities such as the number 2, however, matching spatiotemporal regions seems like the wrong thing to do.

So Lewis's recipe needs to be modified. In what follows I propose one such modification. The modification proceeds via the notion of *grounding* (see Bennett (2011); deRosset (2013); Rosen (2010); Fine (2012)). First, I'll first say a bit about grounding, then I'll outline the modification. Following Schaffer (2009) I take grounding to be a transitive, asymmetric, irreflexive, hyperintensional relation in virtue of which one entity depends upon another. I recognise, of course, that the formal properties of grounding are controversial: Schaffer (2012) has raised counterexamples to transitivity; Jenkins (2011) has offered counterexamples to irreflexivity and Thompson (2016) has outlined counterexamples to asymmetry. For now, however, I will work with the above 'standard package' of grounding features. I will also assume that an entity x *totally* grounds an entity y when x is sufficient on its own to ground y and an entity x *partially* grounds an entity y when x is not sufficient on its own to ground y but there is some plurality of which x is a member that is sufficient to ground y. Total grounding implies necessitation, in at least the following minimal sense: if an entity e grounds an entity $e*$ then, necessarily, if e exists, then $e*$ exists. Grounding can be used to define up the twin notions of fundamental and derivative, to wit:

> x is fundamental $=_{df}$ nothing grounds x.
>
> x is derivative $=_{df}$ something grounds x.

These two notions are exhaustive: every entity is either fundamental or derivative and so grounding induces an ordering over the existing things. Instances of ontological dependence are assumed to be cases of grounding: in every case of ontological dependence where the existence of an entity $e*$ depends upon the existence of an entity e, e grounds $e*$.

Following Schaffer (2016) and Wilsch (2015) I assume that grounding is not chaotic; it is rule-governed, in two respects. First, there are *grounding regularities*. One important type of grounding regularity is the regular grounding of one type of entity by another. Examples of such regularities include: the fact that every number grounds a singleton set, the fact that every truthmaker grounds a truth and the fact that every part grounds a whole. Second, there are *grounding constraints*. One important type of grounding constraint is a constraint on what kinds of entities something may be grounded in. Examples of such constraints include: the fact that singleton sets containing numbers are only grounded in the numbers in question and so cannot be grounded in my shoes; the fact that truths about bees cannot be grounded in facts about hedgehogs; and the fact that a whole may only be grounded in its parts, and so cannot be grounded in something else entirely.

Let us call statements of grounding regularities and grounding constraints: grounding laws. Using the concept of a grounding law, it is possible to define a *grounding miracle* as follows:

> A grounding miracle at ω_m, relative to ω_n, is a violation at ω_m of the grounding laws of ω_n, which are at best the almost-laws of ω_m. The grounding laws of ω_m itself, if such there be, do not enter into it.

By a 'violation' of the grounding laws, I just mean:

> ω_m violates the grounding laws of $\omega_n =_{df}$ There is some statement L that is a grounding law at ω_n but that is false at ω_m.

What are grounding laws, metaphysically speaking? I have no idea, and I have no intention of defending a theory of such things in this paper. But I also don't believe that any deep theoretical understanding of grounding laws is required before they may be put to work. Rather, following Wilsch, we can gain an intuitive grip on the concept of a grounding law by way of analogy with physical laws of nature. Here's Wilsch (2015: 3294) (who calls grounding laws 'laws of metaphysics'):

> Laws of metaphysics are akin to laws of nature in the sense that they guide the development of the world along a dimension. Whereas the natural laws

> work along the temporal dimension, the metaphysical laws work along the axis of fundamentality: from the truths of fundamental physics via the truths of chemistry, biology, and so on, all the way up.

There are two important parallels between the grounding laws and the physical laws of nature. First, physical laws define a modality, namely: the modality of physical necessity. So do the grounding laws. Exactly what that modality is, is up for debate. One natural thought, however, is that the grounding laws define the modality of metaphysical necessity either on their own, or in conjunction with a broader class of metaphysical laws. Second, some physical laws of nature are causal laws: they describe how one event brings about another event over time. Grounding laws are like causal laws: they describe how one entity gives rise to another along the dimension of fundamentality (as Wilsch puts it). Like causal laws, then, the grounding laws are conditional in nature: they determine how one entity e gives rise to another $e*$, if e exists.

If the reader yearns for a more detailed account of the grounding laws, then the one that Wilsch (2015; 2016) provides will suffice. According to this account, there are two kinds of grounding laws: ontological principles and linking principles. Wilsch (2015: 3301) describes the difference thus (a 'constructed entity' is a derivative entity):

> Ontological principles determine which collections of entities give rise to constructed entities by means of particular construction operations. Linking-principles determine which of the constructed objects and properties 'go together' to form facts. The two kinds of principles thus work as a team: ontological principles determine the derivative ontology, and linking-principles determine the derivative facts.

Ontological and linking principles in Wilsch's parlance are grounding regularities in mine. It makes sense to expand Wilsch's account to include what I have called grounding constraints as well. Translated into Wilsch's terminology, we might call these 'constraint principles'. Ontological constraint principles determine what it is that a certain type of entity or fact may *not* be constructed by.

To gain a feel for grounding laws, consider again the number 2 and its singleton, and consider each of the following specific instances of grounding:

1. If the number 1 exists, then the number 1 grounds the singleton $\{1\}$.
2. If the number 2 exists, then the number 2 grounds the singleton $\{2\}$.
3. If the number 3 exists, then the number 3 grounds the singleton $\{3\}$.
4. If the number 4 exists, then the number 4 grounds the singleton $\{4\}$.

...

5. If the number n exists, then the number n grounds the singleton {n}.

The grounding facts about urelements include all of these particular instances of grounding, plus the fact that nothing else grounds a singleton set other than its urelement and whatever grounds that urelement might possess. A basic grounding law for sets and their urelements may therefore be stated as follows:

> (GL_1) For any A, if A exists, then A and the grounds for A (if it has any) ground the existence of {A} and *nothing else does*.

(GL_1) combines a grounding regularity with a grounding constraint: the law dictates what it is that, together, A and the grounds for A ground, and it also places strong constraints on what it is that singleton sets may be grounded in.

Note that (GL_1) implies at least the following four claims (i) for any A, if A exists, then something grounds something; (ii) if A exists then the grounds for A ground something; (iii) every number grounds a singleton set and (iv) every singleton set grounds a singleton set. However, we should not count these more specific claims – what we might call mere grounding regularities – as grounding laws.

In order to sort the grounding laws from the mere grounding regularities that the laws imply we can take a best systems approach. According to the best systems approach to physical laws, the physical laws of nature are the axioms of the system that best unifies our beliefs regarding empirical reality. Unification is understood in terms of a trade-off between strength, understood as deductive power, and simplicity.

A best systems approach to the grounding laws can be sketched as follows. First, we consider all of our beliefs regarding ontological dependence, roughly: beliefs about what ontologically depends on what. Next, we consider the grounding laws to be the axioms of the system that best unifies those beliefs. As before, unification is a trade-off between strength, understood in terms of deductive power, and simplicity.

The best systems approach selects (GL_1) as the law over the mere grounding regularities that it implies. Each of the regularities that (GL_1) implies are more specific than (GL_1) itself, and so imply less. Moreover, these regularities are, at best, as simple as (GL_1) itself. So (GL_1) is superior to these other regularities as a basis for the best system and thus (GL_1) alone will be an axiom. The best systems approach also selects (GL_1) over any conjunction of (GL_1) with another law. For in this case, there is a loss of simplicity with no corresponding gain in deductive power.

There is more to say about grounding and the grounding laws, but the basic idea is, I hope, clear enough: (i) grounding is a relation between entities; (ii) ontological dependence involves grounding; (iii) grounding is lawful, in this sense: there are grounding regularities and grounding constraints and (iv) the difference between grounding laws and mere grounding regularities can be handled via a best systems approach.

Having introduced the concept of grounding, I will now put it to work. Using the concept of grounding, a recipe for determining the closeness of worlds may be stated as follows:

1. It is of the first importance to avoid grounding miracles.
2. It is of the second importance to maximise perfect match with respect to the fundamentals.
3. It is of the third importance to maximise perfect match with respect to the derivatives.

To cast a double slogan: no change to the fundamentals is worth a change to the derivatives and no change to the laws is worth a change to the fundamentals. In the next section I will demonstrate the recipe on a toy model. In demonstrating the recipe I will also explain why it is that avoiding grounding miracles is more important than matching for the fundamentals and why that, in turn, is more important than maximising perfect match with respect to the derivatives.

The toy model is a simple, abstract model of a certain kind of structure: a structure involving entities that depend upon other entities in chains and in a lawful way. The model is thus structurally isomorphic to a range of realistic cases of ontological dependence. Indeed, the toy model mirrors the case of 2 and its singleton almost exactly, and so what I say for the toy model goes for that case as well. On the reasonable assumption that the actual world features well-founded chains of ontological dependence, the success of the account within the toy model provides evidence for the success of the account more generally.

Before turning to the toy model, it is important to offer a couple of points of clarification about the metric. First, with respect to the third line, there are two ways in which the derivatives might differ between a world ω_n and a world ω_m. On the one hand, it may be that there is some derivative entity in ω_n that does not exist in ω_m or *vice versa*. On the other hand, it may be that there is some entity in ω_n that is derivative, and that exists in ω_m (but that is not derivative in ω_m) or *vice versa*. Matching with respect to the derivatives involves differences of the first kind only. Differences of the second kind, I count as differences in the fundamentals.

Second, it is important to emphasise two differences between my account and Lewis's. On Lewis's account, the evaluation of a counterfactual at the actual world is broadly insensitive to what the laws are at other worlds. This shows up in his account of miracles. A world ω features a miracle with respect to the actual world, when something that is a law actually is not a law in ω. Beyond this, however, exactly what the laws are in ω doesn't matter for the closeness of ω. I have used a notion of a miracle that is analogous to Lewis's. However, unlike Lewis's account it does matter what the grounding laws are at non-actual worlds when it comes to determining their closeness. That's because in order to determine the closeness of a world we must compare the worlds in terms of the fundamentals. But presumably, we can only know what is fundamental at a world if we know what the grounding laws are at that world, since it is the grounding laws that tell us what grounds what.

It is useful to spell out this difference a bit further. Consider the closest world in which raising my hand at time t violates an actual law, and in which the past up until t is the same as it is actually. Suppose that, after the hand-raising, the world evolves in accordance with the actual laws. We might well speculate that the laws in this world are the same as the actual laws. But, for Lewis, this is mere speculation and makes no difference to the truth of a counterfactual about my hand raising. On my account, by contrast, we must hold fixed as much as we can about the actual grounding laws in the closest world. This is not to say that differences in grounding laws between worlds are automatically counted as miracles. The idea, rather, is that, setting aside whatever grounding miracles there might be, we need to hold fixed the grounding laws in order to compare worlds with respect to the fundamental and derivative entities in lines 2 and 3 of the recipe.

The second difference between Lewis's account and my own has to do with what it takes to violate a grounding law. For Lewis, a law of nature at a world ω_1, say that all Fs are Gs, is violated at a world ω_2 just in case there is an F at ω_2 that is not a G. The grounding laws are not necessarily violated in the same way. Consider a grounding law that connects the number 2 to the singleton $\{2\}$. There might well be an impossible world in which both 2 and $\{2\}$ exist, and yet the grounding law in question is violated because there is no grounding relation between the two entities.

Is there a problem for my account lurking in these differences with Lewis's view? One might think so, at least with respect to the first difference. When comparing worlds for the purpose of evaluating a counterfactual of ontological dependence, we must consider only those impossible worlds in which there are grounding laws and in which there is grounding. Otherwise, there is little sense to be made of the metric. We cannot compare worlds with respect to fundamentals and derivatives if a world lacks grounding, given how 'fundamental' and 'derivative' are being defined. But then one might worry that, when evaluating the relevant counterfactuals, I am holding fixed facts about grounding in a troubling manner.

To see the worry more clearly, consider that in Lewis's metric he makes reference to *spacetime*. One might raise a similar worry: isn't he holding fixed spatiotemporal relations when proposing this metric and isn't this a problem? But, of course, for Lewis *all* worlds are spatiotemporal, that is how they are defined. So these relations need to be held fixed, but for a principled reason: because it is a background constraint on Lewis's theorising that all of the worlds share the same spatiotemporal relations. By contrast, it is not the case that impossible worlds are defined in terms of grounding. So we cannot take the features of grounding to be a background constraint on theorising in the same way. We need some *further* reason to hold these features fixed.

We also need to avoid holding fixed these features for the wrong reasons. Here's an example of the wrong reason to hold the relevant features fixed: features of grounding are metaphysically necessary and we should hold fixed all metaphysically necessary facts when evaluating counterfactuals of ontological dependence. That is the wrong reason because, when considering impossible worlds, there is no presumption in favour of metaphysically necessary facts.

Of course, one could just take features of the grounding relation like asymmetry to be encoded by the grounding laws. The requirement to hold these features fixed would thus be captured by the first line of the metric. But I don't want to make this assumption. There is nonetheless a good reason to hold those features fixed. Doing so is needed to ensure that the account yields the right truth-values for counterfactuals of ontological dependence. For instance, consider the asymmetry of grounding. Now, take the number 2 and its singleton $\{2\}$ and consider a world in which 2 does not exist but in which $\{2\}$ and $\{\{2\}\}$ symmetrically ground each other. Such a world may be very close to the actual world, given that it differs very little in terms of the fundamentals (only 2 is missing, assuming it is fundamental) and, arguably, it doesn't violate any grounding laws (though this depends a bit on whether the grounding laws are asymmetric). So it looks to be a contender to be the closest possible world, which has the potential to undermine the truth of (CP_1). This outcome is avoided if we ensure that the closest impossible worlds to the actual world are all maximally similar with respect to grounding itself. To do this, we should hold fixed that grounding is asymmetric; not because this is a necessary truth about grounding, but because it pre-selects the right set of impossible worlds. This is no different to other dimensions of the similarity relation which are reverse-engineered in the same way.

As noted, the requirement to hold fixed something about grounding is encoded by the second line of the metric. It could be argued, however, that the second line actually under-specifies how much of grounding is to be held fixed. While we must hold fixed that 'fundamental' and 'derivative' are defined in terms of grounding, perhaps we can do that without holding fixed features of grounding like asymmetry. One way to handle this issue is to add a further line to the metric, one that demands similarity in terms of grounding itself, which is then given primary importance. For my part, however, I am inclined to shift the requirement to hold grounding fixed into the context of evaluation. Thus, I will assume in what follows that in ordinary contexts in which a counterfactual of ontological dependence is evaluated, one holds fixed the relevant features of grounding (such as asymmetry) in order to conduct the evaluation.

5. A Toy Model

This section proceeds in three stages. First, I will outline the toy model (§5.1). Following that I will show how to evaluate a 'would' counterfactual of ontological dependence against the toy model (§5.2). Finally, I will use the toy model to explain why it is that the recipe outlined above requires matching for grounding laws first, fundamentals second and derivatives third (§5.3).

5.1. ω_0

Consider the following world ω_0. In ω_0, there exist four entities: a, b, c and d, which correspond to the entity types A, B, C and D respectively. The world is governed by the following grounding laws:

1. If A's exist, then A's ground B's and C's.
2. If B's exist, then B's ground C's.

3. Only A's may ground B's.
4. Only A's and B's may ground C's.

The grounding structure of the situation is as follows:

1. a grounds b.
2. b grounds c.
3. a is not grounded in anything.
4. d is not grounded in anything.

The world ω_0 may be modeled as follows:

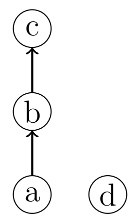

Figure 1

ω_0: arrows represent grounding relations, with the direction of the arrow corresponding to the direction of grounding. Labeled nodes represent entities.

In ω_0, c's existence depends on b's because b grounds c. So, consider the following 'would' counterfactual of ontological dependence:

(CP$_2$) If b had not existed, c would not have existed.

(CP$_2$) ought to be true in ω_0 world: removing b from ω_0 world should force the loss of c as well. If we use the simple account as the basis for evaluating (CP$_2$), we get the wrong result. Recall that, according to the simple account, the closest worlds to ω_0 are the worlds with the most similar ontology. Given the simple account, a world in which b does not exist, but c exists anyway is closer to ω_0 than a world in which b does not exist and c fails to exist as well. That's because there are two ontological differences between ω_0 and a world in which b does not exist and c does not exist (the lack of b and the lack of c), whereas there is only one difference between ω_0 and a world in which b does not exist and c exists anyway (the lack of b). The reasoning here reflects the Nixon case outlined in the previous section exactly.

Analysis 2 in combination with my three step recipe for ordering worlds vindicates (CP$_2$). The evaluation of that counterfactual proceeds as follows. First, we take the set of worlds in which neither b nor c exist. Call such worlds the $\mathbb{W}_1$ worlds (henceforth I will use '$\mathbb{W}$' to refer to types of worlds, reserving 'ω' for specific worlds). Next we identify the set of worlds in which b does not exist but c exists anyway. This set of worlds divides into two broad kinds, depending on how c is recovered. First,

c may be recovered in the absence of b by being added as a new fundamental entity. Call these the $\mathbb{W}_2$ worlds. Second, c may be recovered in the absence of b by being added as a derivative entity, one that is grounded in something other than b. Call these the $\mathbb{W}_3$ worlds. Since there is no alternative way to recover c in the absence of b, there are no other worlds to consider when evaluating (CP_2). Accordingly, if there is some $\mathbb{W}_1$ world that is closer than any $\mathbb{W}_2$ or $\mathbb{W}_3$ world, then (CP_2) is true, and if not, then not.

In order to proceed further with the evaluation of (CP_2), I need to say a bit more about each of the worlds just specified. With respect to $\mathbb{W}_1$ worlds, $\mathbb{W}_2$ and $\mathbb{W}_3$ worlds, all three types of world divide into two further kinds, depending on how b has been removed. On the one hand, b may be removed by removing a – the fundamental entity that grounds it – thereby excising the entire grounding chain that gives rise to b. On the other hand, b may be removed by keeping a in place and violating the grounding law that if A's exist, then A's ground B's and C's. Let us call those $\mathbb{W}_1$, $\mathbb{W}_2$ and $\mathbb{W}_3$ worlds in which b has been removed by removing a the $\mathbb{W}_{1.1}$, $\mathbb{W}_{2.1}$ and $\mathbb{W}_{3.1}$ worlds. And let us call those $\mathbb{W}_1$, $\mathbb{W}_2$ and $\mathbb{W}_3$ worlds in which b has been removed by keeping a in place and violating the grounding law that if A's exist, then A's ground B's and C's, the $\mathbb{W}_{1.2}$, $\mathbb{W}_{2.2}$ and $\mathbb{W}_{3.2}$ worlds.

With respect to the $\mathbb{W}_3$ worlds, there are two ways to re-introduce c as a derivative entity: ground it in something that exists in ω_0 and potentially break a grounding law (namely the law that forbids the grounding of c in something other than A's and B's) or ground it some new entity, e. I won't consider worlds in which c is grounded in some new entity e here. Grounding c in e will require breaking a grounding law, so such worlds are at least as far away as worlds in which c is grounded in something that exists in ω_0 (something that does not ground c in that world). However, worlds in which c is grounded in some new entity e will also require either a new fundamental (if e is fundamental) or a new derivative (if e is derivative). Adding e and using it to ground c does not *reduce* the number of differences between $\mathbb{W}_3$ worlds and ω_0.

Having specified each world more fully, let us now suppose that $\mathbb{W}_1$, $\mathbb{W}_2$ and $\mathbb{W}_3$ worlds differ only from ω_0 in so far as they must in order to possess the features outlined above. To get a feel for what each world looks like, it is useful to catalogue the differences between each world specified and ω_0. In doing so, I will count differences between worlds as follows.

1. A world ω_m differs from a world ω_n with respect to the grounding laws, when ω_m possesses a grounding miracle with respect to ω_n or *vice versa*.

2. A world ω_m differs from a world ω_n with respect to the fundamentals, when the set of fundamental entities at ω_m differs from the set of fundamental entities at ω_n.

3. A world ω_m differs from a world ω_n with respect to the derivatives, when there is some entity that exists at ω_m and is derivative in that world that does not exist at ω_n or *vice versa*.

Now, consider the $\mathbb{W}_1$ worlds. In $\mathbb{W}_{1.1}$ worlds, a, b and c do not exist. Compared to $\mathbb{W}_0$ worlds, then, $\mathbb{W}_{1.1}$ worlds possess two differences with respect to the derivatives, one difference with respect to the fundamentals and no differences with respect to the grounding laws, since none are violated in order to remove b. In $\mathbb{W}_{1.2}$ worlds, by contrast, a exists while b and c do not. Compared to ω_0, then, $\mathbb{W}_{1.2}$ worlds possess two differences with respect to the derivatives, no differences with respect to the fundamentals and one difference with respect to the grounding laws, since a law must be violated in order to keep a and remove b. $\mathbb{W}_{1.1}$ and $\mathbb{W}_{1.2}$ worlds are represented in Figure 2.

Figure 2

$\mathbb{W}_1$ *worlds.*

Consider, next, the $\mathbb{W}_2$ worlds. In $\mathbb{W}_{2.1}$ worlds, c exists and is fundamental and neither a nor b exist. So compared to ω_0, $\mathbb{W}_{2.1}$ worlds possess two differences with respect to the fundamentals, one difference with respect to the derivatives and no differences with respect to the grounding laws, since none are violated in order to remove b.

In $\mathbb{W}_{2.2}$ worlds, by contrast, c exists and is fundamental, a exists and b does not exist. So compared to ω_0, $\mathbb{W}_{2.2}$ worlds possess one difference with respect to the fundamentals, one differences with respect to the derivatives and one difference with respect to the grounding laws, since a law must be violated in order to keep a and remove b. The $\mathbb{W}_2$ worlds are represented in Figure 3.

Figure 3

$\mathbb{W}_2$ *worlds.*

Consider, finally, the $\mathbb{W}_3$ worlds. In $\mathbb{W}_{3.1}$ worlds, a and b don't exist, but c exists and is grounded in something that exists in $\mathbb{W}_0$ worlds. Since only d is available for grounding c, that work must be done by d. Compared to ω_0, then, $\mathbb{W}_{3.1}$ worlds possess one difference with respect to the fundamentals, one difference with respect to the derivatives and one difference with respect to the grounding laws, since if d grounds c, then the law 'only A's and B's may ground C's' must be violated in order to allow d to do this grounding work.

In $\mathbb{W}_{3.2}$ worlds, a exists but b does not exist, and c is grounded in something that exists in ω_0. The smallest change we can make in order to ground c is to ground it in a, since no grounding miracles are required to do so (larger changes will only make the world further away from ω_0 and it is the closest worlds we want to consider). Compared to ω_0, then, $\mathbb{W}_{3.2}$ worlds possess no differences with respect to the fundamentals, one difference with respect to the derivatives and one difference with respect to the grounding laws, since the law 'if A's exist, then A's ground B's and C's' must be violated.

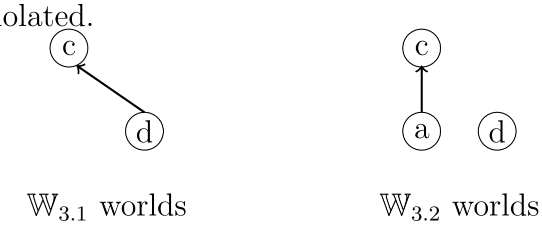

Figure 4

$\mathbb{W}_3$ *worlds.*

The differences between $\mathbb{W}_1$, $\mathbb{W}_2$, $\mathbb{W}_3$ worlds and ω_0 are summarised in Table 1.

World	Grounding Miracles	Differences in Fundamentals	Differences in Derivatives
ω_0	0	0	0
$\mathbb{W}_{1.1}$	0	1	2
$\mathbb{W}_{1.2}$	1	0	2
$\mathbb{W}_{2.1}$	0	2	1
$\mathbb{W}_{2.2}$	1	1	1
$\mathbb{W}_{3.1}$	1	1	1
$\mathbb{W}_{3.2}$	1	0	1

Table 1: Worlds for evaluating (CP$_2$) in $\mathbb{W}_0$.

5.2. 'if b had not existed, c would not have existed'

We are now in a position to show that (CP$_2$) is true. According to **Analysis 2**, (CP$_2$) is true when some world in which b does not exist and c does not exist is closer than any world in which b does not exist and c does exist. If we apply my three-step recipe for determining closeness between worlds, then the list of worlds in Table 1 may be re-ordered in terms of closeness to ω_0 as follows:

World	Grounding Miracles	Differences in Fundamentals	Differences in Derivatives	Rank
ω_0	0	0	0	1st
$\mathbb{W}_{1.1}$	0	1	2	2nd
$\mathbb{W}_{2.1}$	0	2	1	3rd
$\mathbb{W}_{3.2}$	1	0	1	4th
$\mathbb{W}_{1.2}$	1	0	2	5th
$\mathbb{W}_{3.1}$	1	1	1	6th
$\mathbb{W}_{2.2}$	1	1	1	7th

Table 2: ω_0, $\mathbb{W}_1$, $\mathbb{W}_2$ and $\mathbb{W}_3$ worlds ordered using a three-step recipe that prioritises matching grounding laws then fundamentals then derivatives.

In Table 2, those worlds that match ω_0 with respect to the grounding laws are the most similar worlds. Because a $\mathbb{W}_{1.1}$ world is closer than any $\mathbb{W}_2$ or $\mathbb{W}_3$ world, it follows that if b had not existed, c would not have existed. The three-step recipe, then, yields the right result for the toy case.

5.3. Alternative Recipes

Above I promised to explain two features of the recipe: first, why it is that matching with respect to the grounding laws is more important than matching with respect to the fundamentals and, second, why it is that matching with respect to the derivatives is of the least importance.

Both features of the recipe can be explained using the toy case. First, however it is important to answer a prior question, namely: why are there any priorities in the recipe at all? The answer to this prior question is straightforward. If there are no priorities in the recipe, then $\mathbb{W}_1$, $\mathbb{W}_2$ and $\mathbb{W}_3$ worlds will be ranked with respect to ω_0 worlds simply in terms of the total number of differences between those worlds and ω_0. When differences in the derivatives, fundamentals and grounding laws are weighed equally, (CP$_2$) is false. That's because there is no world in which b does not exist and c does not exist that is closer than any world in which b does not exist and c does exist; a $\mathbb{W}_{3.2}$ world is the closest world to ω_0. To see this, simply re-order Table 1 to weigh these three aspects equally. The point generalises beyond the toy case: if fundamentals, derivatives and grounding miracles have equal weight, those worlds that match actuality with respect to the derivatives will tend to dominate all other worlds for closeness. In light of the Nixon case from §3, this is precisely what we want to avoid. So a system of priorities is needed.

Given that there needs to be some prioritisation between differences in the fundamentals, grounding laws and derivatives why prioritise the grounding laws over the fundamentals? To answer this question, suppose we prioritise the fundamentals over the grounding laws. This alternative recipe gets the wrong results in the toy case. That's because, once again, there is no world in which b does not exist and c does not exist that is closer than any world in which b does not exist and c does exist; a $\mathbb{W}_{3.2}$ world is the closest world to ω_0. Again, to see this simply re-order Table 1 so that matching with respect to the fundamentals is prioritised. Such a re-ordering ranks a $\mathbb{W}_{3.2}$ world more highly than any other world. So matching the grounding laws should be more important than matching the fundamentals.

This brings us to the second of the two questions posed above: why prioritise matching the grounding laws and the fundamentals over matching the derivatives? There are two ways to alter the priority weighting so that matching the derivatives is more highly valued. We could either value derivatives more highly than any other feature, or we could allow that the grounding laws are of primary importance and then value derivatives over fundamentals. Both recipes yield the wrong results. When the derivatives are prioritised above all else, either a $\mathbb{W}_{3.2}$ world or a $\mathbb{W}_{2.1}$ world is the closest world to ω_0 (depending on whether fundamentals are weighed more heavily than the grounding laws or vice versa). When the derivatives are prioritised over the fundamentals, a $\mathbb{W}_{2.1}$ world is the closest world to ω_0. In both cases (CP_2) is false. Because there is no other way to value the derivatives over the fundamentals, it follows that matching the fundamentals should be more important than matching the derivatives.

6. Conclusion

My three-step recipe – when combined with **Analysis 2** – correctly yields the result that (CP_2) is true. It can be shown that the account yields the right result in the example of 2 and {2}, in essentially the same way (I will forego the details). The account is also immune to Nixon-style counterexamples. Nixon cases for 'would' counterfactuals of ontological dependence all have a similar structure. As noted, a 'would' counterfactual of ontological dependence is a counterfactual with the following broad form: if a had not existed, b would not have existed. A Nixon-style problem for such counterfactuals can therefore be constructed as follows. First, identify a range of entities $c_1 ... c_n$ that depend for their existence on b in the actual world (if there are any such entities). Second, consider two kinds of world: (i) a world in which a does not exist, b does not exist and the c_n do not exist (a $\mathbb{W}_1$ world), and (ii) a world in which a does not exist and b does exist, along with the c_n (a $\mathbb{W}_2$ world). Third, note that a $\mathbb{W}_2$ world is more similar to actuality than any $\mathbb{W}_1$ world because the c_n exist in $\mathbb{W}_2$ and not in $\mathbb{W}_1$. Finally, conclude that 'if a had not existed, b would not have existed' is false because there is no world in which a does not exist and b does not exist that is closer than any world in which a does not exist and b does exist.

According to my three-step recipe a $\mathbb{W}_2$ world in the construction just specified will always be further away from actuality than a $\mathbb{W}_1$ world in that construction and so the case fails at the third step. That's because a further change must be made to a $\mathbb{W}_2$ world in order to reintroduce b and, with it, the c_n in the absence of a. Either some grounding law will need to be broken, because b is grounded in something that does not ground it actually in a $\mathbb{W}_2$ world, or there will need to be some difference in the fundamentals, because b is fundamental in a $\mathbb{W}_2$ world and not actually. Either way, these differences will outweigh the similarity bought back by recovering the c_n, rendering the 'would' counterfactual 'if a had not existed, b would not have existed' true.

One limitation of the toy model is that it features only well-founded chains of ontological dependence, and so does not demonstrate the applicability of the recipe to ontological structures that feature infinite descent. Moreover, it is unclear that the recipe is applicable to these cases in principle, given that it relies on the priority of fundamentals. However, even if, ultimately, nothing is fundamental, this does not render the account of the similarity relation that I have provided useless. For even if nothing is actually fundamental, one can assume within a context that something is fundamental by treating a particular part of a chain of ontological dependence as the terminus. One can then contextually restrict interest to laws that govern just that fragment of the ontological chain one has isolated. In this way, the evaluation of 'would' counterfactuals can proceed as in the toy model.

Still, it would be nice to have a recipe that works equally well for cases of infinite descent and cases in which something is fundamental. But perhaps that is just too much to ask. It may turn out that we need different accounts of the similarity relation to handle these two types of cases. I leave this as an open question. For now, it is enough to note that the recipe works for a particular type of ontological structure and that it is a common enough view that this type of structure is widespread. The upshot being that the recipe will work for a number of 'would' counterfactuals of ontological dependence.

I recognise, of course, that there may be other recipes that would work just as well. That being said, the present recipe is quite a natural one to use. For the recipe appeals to facts about grounding and fundamentality that are important to understanding ontological dependence more generally. When we evaluate counterfactuals of ontological dependence in this way, we are therefore evaluating them in terms that appear quite relevant to the counterfactuals at issue. At any rate, I welcome the suggestion of alternative recipes. Thinking through each such recipe helps us to gain a better understanding of 'would' counterfactuals of ontological dependence. Gaining a better understanding of these 'would' counterfactuals is important. A great deal of reasoning within philosophy makes use of these counterfactuals in one way or another.

Sam Baron
University of Western Australia
samuel.baron@uwa.edu.au

Acknowledgements I would like to thank Ray Briggs, Jonathan Schaffer and Alastair Wilson for comments on earlier drafts of this paper. I would also like to think the attendees at a Rutgers postgraduate metaphysics reading group where a very early version of this paper was presented. Finally, I would like to thank two referees of this journal for their extremely helpful feedback. Work on this paper was partly funded by an Australian Research Councial Discovery Early Career Researcher Award (DE180100414).

References

Beall, Jc, and Bas C. van Fraassen. (2003) *Possibilities and Paradox: An Introduction to Modal and Many-valued Logic*. Oxford: Oxford University Press.

Bennett, Karen. (2011) 'By our bootstraps.' *Philosophical Perspectives*, 25: 327–353.

Berto, Francesco, Rohan French, Graham Priest, and David Ripley. (2018) 'Williamson on counterpossibles.' *Journal of Philosophical Logic*, 47(4): 693–713.

Bjerring, Jens Christian. (2014) 'On Counterpossibles' *Philosophical Studies*, 168(2): 327–353.

deRosset, Louis. (2013) 'Grounding explanations.' *Philosophers' Imprint*, 13: 1–26.

Fine, Kit. (1975) 'Review of "counterfactuals".' *Mind*, 84: 451–458.

Fine, Kit. (2012) 'A guide to ground.' In *Metaphysical Grounding: Understanding the Structure of Reality*. Edited by Fabrice Correia and Benjamin Schneider, 37–80. Cambridge: Cambridge University Press.

Jago, Mark. (2015) 'Impossible worlds.' *Noûs*, 49: 713–728.

Jenkins, Carrie I. (2011) 'Is metaphysical dependence irreflexive?' *The Monst*, 94: 267–276.

Kim, Jaegwon. (1973) 'Causes and Counterfactuals' *Journal of Philosophy*, 70(17): 570–572.

Kment, Boris. (2014) *Modality and Explanatory Reasoning*. Oxford: Oxford University Press.

Laan, David Vander. (2004) 'Counterpossibles and similarity.' In *Lewisian Themes: The Philosophy of David K. Lewis*. Edited by Frank Jackson and Graham Priest, 258–276. Oxford: Oxford University Press.

Lewis, David. (1973) 'Counterfactuals and comparative possibility.' *Journal of Philosophical Logic*, 2: 418–446.

Lewis, David. (1979) 'Counterfactual dependence and time's arrow.' *Noûs*, 13: 455–476.

Lewis, David. (1987) 'Events' *Philosophical Papers II*, 242–270. Oxford: Oxford University Press.

Mares, Edwin D. (1997) 'Who's afraid of impossible worlds?' *Notre Dame Journal of Formal Logic*, 38: 516–526.

Nolan, Daniel. (1997) 'Impossible worlds: A modest approach.' *Notre Dame Journal of Formal Logic*, 38: 535–572.

Priest, Graham. (2002) *Beyond the Limits of Thought*. Oxford: Oxford University Press.

Restall, Greg. (1997) 'Ways things can't be.' *Notre Dame Journal of Formal Logic*, 38: 583–596.

Ripley, David. (2012) 'Structures and circumstances.' *Synthese*, 189: 97–118.

Rosen, Gideon. (2010) 'Metaphysical dependence: Grounding and reduction.' In *Modality: Metaphysics, Logic and Epistemology*. Edited by Bob Hale and Aviv Hoffman, 109–135. Oxford: Oxford University Press.

Schaffer, Jonathan. (2009) 'On what grounds what.' In *Metametaphysics*. Edited by David Manley, David J. Chalmers, and Ryan Wasserman, 347–383. Oxford: Oxford University Press.

Schaffer, Jonathan. (2012) 'Grounding, transitivity and contrastivity.' In *Grounding and Explanation*. Edited by Fabrice Correia and Benjamin Schneider, 122–138. Cambridge: Cambridge University Press.

Schaffer, Jonathan. (2016) 'Grounding in the image of causation.' *Philosophical Studies*, 173: 49–100.

Thompson, Naomi. (2016) 'Metaphysical interdependence.' In *Reality Making*. Edited by Mark Jago, 38–56. Oxford: Oxford University Press.

Williamson, Timothy. (2013) *Modal Logic as Metaphysics*. Oxford: Oxford University Press.

Wilsch, Tobias. (2015) 'The nomological account of ground.' *Philosophical Studies*, 172: 3293–3312.

Wilsch, Tobias. (2016) 'The deductive-nomological account of metaphysical explanation.' *Australasian Journal of Philosophy*, 94: 1–23.