Lewis' Account of Counterfactuals is Incongruent with Lewis' Account of Laws of Nature

Foad Dizadji-Bahmani and Seamus Bradley
DRAFT VERSION OF PAPER PRESENTED AT BSPS2014

July 14, 2014

1 Introduction

In this paper we argue that there is a problem with the conjunction of David Lewis' account of counterfactual conditionals and his account of laws of nature. This is a pressing problem since both accounts are individually plausible, and popular.

There is a well-known objection to Lewis' account of counterfactuals, the most famous instance of which is the so-called "nuclear button" example due to Fine (1975). In response to this objection Lewis clarified his original account of counterfactuals. What we show is that Lewis' modified account is incongruent with his account of laws of nature.

We proceed as follows. We start by setting out Lewis' account of counterfactuals. We then present the objection and Lewis' response thereto. We then set out Lewis' account of laws of nature. Finally we show in what sense the two accounts, of counterfactuals and of laws of nature, are incongruent.

2 Counterfactuals

You are holding a ball in front of you. You are considering letting go of it. As a matter of fact you do not let go of the ball: you hold it stationary, arm outstretched. Now, consider the following two counterfactuals:

(i) If you had let go of this ball, it would have dropped to the floor.

(ii) If you had let go of this ball, it would have turned into a badger.

The first is true and the second is false. Why? Both have the same logical form so it must be a difference in the semantics that accounts for the difference in truth value. Lewis famously gave an account of the semantics of counterfactual conditionals in terms of possible worlds, and, more specifically, the relative similarity of various possible worlds to the actual world (Lewis, 1973).[1]

Here is what his theory says about our examples. Consider (i) first. Let's formalise this sentence as "$L \boxright F$". Consider the possible worlds where you did in fact let go of the ball; worlds where L is true. Call these 'L-worlds'.[2]

And consider also possible worlds where the ball drops to floor; worlds where F is true. Call these 'F-worlds'. The truth-value of (i) depends on the relative similarity – the *closeness* – of different L-worlds to the actual world, as follows: (i) is true iff there is a $L \wedge F$-world closer to the actual world than any $L \wedge \neg F$-world.

Now consider (ii). Formalise it as '$L \boxright B$', and consider again L-worlds but this time also B-worlds, worlds where the ball turns into a badger. Once again the truth-value of the counterfactual depends on the relative similarity – the *closeness* – of different L-worlds to the actual world but this time it is the $L \wedge B$-worlds and $L \wedge \neg B$-worlds which are salient: (ii) is true iff there is a $L \wedge B$-world closer to the actual world than any $L \wedge \neg B$-world.

Intuition deems (i) true because we think that a world in which the ball is let go and drops is a world similar to our own, while a world where the ball does not drop when let go is somehow odd, deviant, far removed. On the other hand intuition deems (ii) false precisely because we think that a world in which the ball is let go of and turns into an badger *is* odd, deviant, far removed. Lewis' account is compelling because it captures these intuitions in a clear and systematic way, or so it is claimed.

For the general case of a counterfactual like "$A \boxright C$" the analysis takes this form:

- $A \boxright C$ is true iff the closest $A \wedge C$-world is closer than all $A \wedge \neg C$-worlds.

This sketch of the analysis obviously leaves many of the details underexplored, but it suffices for the present purpose of setting up the problem with which we are concerned. The important point for the coming discussion is that the truth-value of counterfactuals turns on a notion of *similarity* – or *closeness* – of possible worlds.

[1] The foundations of this account of the semantics of counterfactuals goes back to Stalnaker (1968), but it is Lewis' elaboration of it that will be of interest to us in what follows.

[2] As is well-known Lewis advocated a 'counterpart' theory of possible worlds. So strictly speaking the possible worlds to consider are not worlds where *you* do something, but, rather, where your *counterpart* does something. (Whether this is the right account is not our concern.) For convenience, we'll gloss over this and we'll speak in the 'identity mode' throughout.

3 A Flaw in the Analysis

It quickly became apparent that there was a flaw with Lewis' analysis of counterfactuals.[3] Let's consider this flaw with respect to our example. Lewis' account made counterfactual (i) true by virtue of the intuitive claim that there is a world in which you let go of the ball and the ball drops which is closer than all other worlds in which you let go and the ball does not drop. That is, worlds where you let go of the ball and it drops are less odd, deviant, or far-removed than worlds where you let go and the ball does not drop. But consider the following specific possible world. It is like the actual world up to the time at which you did *not* let go of the ball. This sounds odd but recall that in the actual world you considered letting go of the ball but didn't. In the possible world under consideration, at the time of deliberation you do the opposite–you *do* let go. But at this possible world, once you let go of it, the ball stays where it was. It just hovers there, exactly *as if* you were still holding it.

There are a lot of possible worlds floating about, so let us take stock. We are interested in the counterfactual (i) formalised $L \mathrel{\Box\!\!\rightarrow} F$: "If you had let go of the ball, it would have dropped to the floor". The following possible worlds are relevant to the analysis.

(a) The actual world: you do not let go of the ball but hold it stationary. This is a $\neg L \wedge \neg F$-world.

(b) The standard world: the world where you let go of the ball and it falls to the ground. This is a $L \wedge F$-world.

(c) The deviant world: the world where you let go of the ball but it stays in the same position as it does in (a). This is an $L \wedge \neg F$-world.

At (c), L is true; (c) is a L-world. But (c) is constructed[4] in such a way as to be like the actual world – (a) – in all other respects. So, for example, once your fingers stretch out and straighten, your arm stays in the same place, and you don't otherwise move. The ball stays just as it was before you let go: it hovers in the same position unchanging, even down to the distribution of pressure on its surface that your hand caused before, and down further still! (c) is exactly like (a) except for whatever needs to be different to make L come out true. As we'll say, (c) is *maximally similar excepting* L to the actual world.

Intuitively, counterfactual (i) is true, and Lewis' account must do justice to this. Under Lewis' account, for counterfactual (i) to be true requires (b) to be closer to the actual world, (a), than any world where I let go and the ball does not drop (any $L \wedge \neg F$-world). That is, under Lewis' account, (b), which is a $L \wedge F$-world, needs to be closer to the actual world, (a), than any $L \wedge \neg F$-world. However, (c) is a $L \wedge \neg F$-world which appears to be closer to (a) than

[3] The most well-known instance is due to Kit Fine (Fine (1975), see also Bennett (1984)). However, the objection is general; we present it in terms of our example.

[4] Of course, if modal realism is true then (c) is not 'constructed' but is, rather, 'sitting out there'. Nothing hangs on this. Modal realists can think of 'constructing' a possible world as merely directing one's attention to a possible world.

(b). Indeed (c) is in a sense *maximally similar excepting L* to the actual world. So it seems like under Lewis' account, counterfactual (i) comes out false, contrary to the intuition to which his account must do justice.

It should be clear that nothing hangs on the specifics of the example. Here's a general recipe for finding such deviant worlds.

- Take any supposedly true counterfactual with a false antecedent $A \Box\!\!\rightarrow C$.
- Take the $A \wedge C$-world which allegedly makes the counterfactual true.[5]
- Concoct a world as follows: "A, but everything else is as if $\neg A$". Call this a deviant world.
- Since $\neg A$ is true in the actual world, it looks like this deviant world *has to be* closer to the actual than the world that allegedly made the counterfactual true.
- In this deviant world, C has whatever truth value it has in the actual world.
- So on Lewis' account all counterfactuals come out as having the same truth values as their consequents have in the actual world.

Given that not all counterfactuals have the same truth-values as their consequents have in the actual world, Lewis' account is seemingly flawed.

4 Similarity and Laws of Nature

Lewis responded to this problem by arguing that the notion of closeness of worlds used in arguing that these deviant worlds are closer to the actual than the respective standard worlds is erroneous. They are closer to the actual world if similarity is measured simply in terms of particular matters of local fact. This is because such possible worlds are *constructed* so as to be maximally similar to the actual world with respect to all those factors that don't impact on the antecedent of the conditional. But, Lewis argues, this is the wrong measure of similarity.

Prompted by Fine's famous "nuclear button" example (cf. Fine (1975)), which is a particularly dramatic instance of the general problem discussed above, Lewis constructed a hierarchy of respects in which worlds can be similar or dissimilar. Lewis' claim is that a similarity relation respecting this hierarchy doesn't make the deviant possible world closer to the actual world than the relevant $A \wedge C$-world that is to make $A \Box\!\!\rightarrow C$ true.

How so? In short, Lewis' hierarchy ties similarity to the laws of nature. The worlds closest to the actual world are those which have the same unviolated

[5]Lewis' analysis isn't committed to their being a single such world. That is, Lewis was not committed to the Uniqueness assumption. We talk in terms of "the" possible world, but everything works just fine if you pick any one of the closest possible worlds.

laws of nature. The difference between these worlds and the actual world are just differences between mere contingent facts, differences which do not involve violation of laws at all. Further are worlds that have the same laws but differ only in some "small miracle" – some minor violation of the laws. Further still are worlds in which there are "big miracles" – some major violations of the laws.[6] And furthest are worlds which have different laws altogether. Or better: worlds where things are so different to the actual world that it stops being useful to characterise these worlds in terms of violations (minor or major) of the laws at the actual world, but which are better characterised as worlds where the laws are different (Lewis, 1979).

The deviant $L \wedge \neg F$-world that seemed to make the intuitively true counterfactual (i) false involves a big miracle to keep the ball hovering there, Lewis would claim. It involves, *inter alia*, the violation of the laws of gravity. The standard $L \wedge F$-world that makes the intuitively true counterfactual true only involves the tiny miracle that you chose to drop the ball rather than not.[7] So, as per the hierarchy, the standard $L \wedge F$-world is closer to the actual world than is the deviant $L \wedge \neg F$-world. Thus, via the hierarchy Lewis recovers the intuitively correct truth-value assignments of counterfactuals.

Clearly laws of nature play a central role in this hierarchy, and, so, play a central role in determining the truth-values of counterfactuals. Our overall claim is that Lewis' conception of laws of nature is incongruent with his account of counterfactuals in the following sense: Lewis needs to take similarity in terms of the laws as the important kind of similarity in order for the theory of counterfactuals to work; however, such a notion of similarity is hard to reconcile with his theory of the laws of nature. To see this we need to examine Lewis' account of laws of nature. We turn to this now.

5 Lewis' Account of Laws of Nature

What is a law of nature? Lewis advocated what is now commonly called the 'best systems' account (BSA) of laws of nature in response to this question.[8] On Lewis' account the laws of nature are generalisations which are the theorems of that true deductive system which strikes the best balance between simplicity and strength. Lewis puts it this way:

[6] Lewis introduced this distinction between small and big miracles, but, of course, this is a matter of degree not kind. We'll say more about this in the section 5. Note also that talk of 'violations' of laws is, strictly speaking, a *façon de parler* because on Lewis' account the laws of nature at a particular world are *true* generalisations (of a particular kind) at that world, and true generalisations cannot be violated. (For that would make them, strictly speaking, false!) This will become clearer in section 5.

[7] This involves a miracle, because we are assuming that the histories of the worlds up until the time of the dropping of the ball are identical and that the world is deterministic.

[8] This is also known as the 'MRL account' (Mill-Ramsey-Lewis) since precursors of the central idea that Lewis systematised can be found in J.S. Mill and Frank Ramsey. BSA is popular amongst contemporary philosophers. Cf. Cohen and Callender (2009) for a recent overview and their modification.

> Take all deductive systems whose theorems are true. Some are simpler, better systematized than others. Some are stronger, more informative, than others. These virtues compete: an uninformative system can be very simple, an unsystematized compendium of miscellaneous information can be very informative. The best system is the one that strikes as good a balance as truth will allow between simplicity and strength. How good a balance that is will depend on how kind nature is. A regularity is a law iff it is a theorem of the best system. (Lewis, 1994, p.478).

How are we to understand *simplicity* and *strength*, and how they are *balanced*? Whilst we take it that the concepts are intuitively clear, it is difficult to give a precise answer to this question. Lewis himself notes that it would be unacceptable for his account to ground *simplicity* and *strength* only by our psychology. But how to ground them otherwise? And how to balance them against each other? Lewis proposes to not settle these questions but to avoid them:

> I suppose our *standards* of simplicity and strength and balance are only partly a matter of psychology. It's not because of how we happen to think that a linear function is simpler than a quartic or a step function; it's not because of how we happen to think that a shorter alternation of prenex quantifiers is simpler than a longer one; and so on. Maybe some of the exchange rates between aspects of simplicity, etc., are a psychological matter, but not just anything goes. If nature is kind, the best system will be robustly best – so far ahead of its rivals that it will come out first under any standards of simplicity and strength and balance. We have no guarantee that nature is kind in this way, but no evidence that it isn't. It's a reasonable hope. (Lewis, 1994, p.479)

For the time being, let us suppose that it is indeed the case that the best system is robustly best, in Lewis' sense. Of course, for Lewis' account to be tenable these central notions would nonetheless need to be made more perspicuous. However, the above characterisation suffices for the purpose of setting up the problem with which we are concerned.

Lewis is also committed to Humean Supervenience (HS): "the thesis that the whole truth about a world like ours supervenes on the spatiotemporal distribution of local qualities" (Lewis, 1994, p.473) Under HS, the truth-values of propositions are ultimately determined by the properties instantiated at various spacetime points (or something similar).[9] In fact, Lewis explicitly says that

[9]HS is a central tenet of Lewis' philosophy. Cohen and Callender have (correctly) stressed that HS is independent of BSA (Cohen and Callender, 2009). But for Lewis, HS *is* part of the overall account of laws which he advocates. Lewis acknowledges that HS is not a necessary truth–there are possible worlds where HS is false–but such worlds are very far from the actual world. Indeed, in essence, Lewis is committed to the following: the truth-values of common/ordinary counterfactuals are determined by the relative distances of *only* HS-worlds to the actual world.

much of his work can be seen as contributing to the systematic philosophical position he dubs Humean supervenience:

> Many of [my] papers. . . seem to me in hindsight to fall into place within a prolonged campaign on behalf of the thesis I call "Humean supervenience". . . Humean supervenience is named in honor of the greater denier of necessary connections. It is the doctrine that all there is to the world is a vast mosaic of local matters of particular fact, just one little thing and then another. . . For short: we have an arrangement of qualities. And that is all. There is no difference without difference in the arrangement of qualities. All else supervenes on that. (Lewis, 1987, pp. ix–x)

On Lewis' view, then, the laws of nature at a world are determined by the particular matters of local fact at that world. So the laws of nature at the worlds relevant to the analysis of a counterfactual are determined by the particular facts at the relevant worlds. For they (the laws) are the theorems or axioms of the true deductive system which best summarizes those facts.

Before moving on, let's say a little bit more about what Lewis means by a "violation" of a law of nature. This idea was introduced in the previous section, but we are now in a position to say a little more about it. Consider a world where it is a law that all Fs are Gs. Now consider a world where there is one F that isn't a G. From the perspective of the former world, the latter world looks like a world with the same laws but with one violation of them. This isn't strictly speaking true, since "All Fs are Gs" is *not* a law of the latter world: it is not a true generalisation. But there's an intuitive sense in which we can understand differences between worlds as violations of the laws of one world in another world. It is in *this* sense that we understand Lewis' talk of "violations" of laws.

6 The Incongruence of Lewis' Accounts

Let's make a distinction between two kinds of closeness of possible worlds: the notion of closeness at the level of the Humean mosaic; and closeness in terms of the laws. Let's call these 'm-closeness' and 'l-closeness'; likewise 'm-similarity' and 'l-similarity'.

Imagine a very simple example of a "Humean mosaic" where every point in space can either have some property (being "on") or not (being "off"), like a monochrome screen displaying a picture. Each point in the mosaic is a "pixel". Presumably, whatever Lewis or Lewisians take the mosaic to be, it's much more complicated than this, but the analogy will be fruitful. One candidate explication of what m-similarity between pictures is, is the "Hamming distance": the number of pixels you would have to switch from on to off or from off to on

---

That is, by the relative distances of, to use Lewis' phrase, *worlds which fall within a shell where HS is true*. As we will see, denying HS is one way resolving the incongruence, but one to which Lewis would not be amenable.

in order to turn one picture into the other. This is something like the kind of distance we appealed to in constructing the deviant possible world: the idea was that the deviant world was different at fewer points in the Humean mosaic than the standard world was, and was thus closer to the actual. Now imagine a picture and its "negative": the picture you get from inverting all of the pixels from on to off and from off to on. We would intuitively take these pictures to be very similar, but Hamming distance doesn't capture this similarity: they are different at every pixel. So, similarity of pictures displayed on a computer screen supervenes on the array of pixels displayed. That is, you can't have a difference in similarity without a difference in pixels displayed. Despite this supervenience, similarity of pictures isn't simply a matter of counting differences on the pixel-level (i.e. the analog of m-closeness). Similarity of pictures involves recognising higher-level features that the pictures share or don't; two pictures may be more similar to each other than a third, even if the first is Hamming-distance closer – closer at the level of pixels – to the third than it is to the second. This response is at first pass intuitively appealing: indeed pictures may be more similar despite Hamming-distance closeness as suggested. But notice that the claim makes (tacit) recourse to a notion of similiarity other than that of Hamming-distance. And when it comes to pictures on a screen, we do have an intuitive idea about what similarity amounts to.

Laws of nature, like pictures on a screen, are built up out of the basic building blocks (properties at spacetime points, pixels). And like in the case of pictures on a computer screen, similarity of the higher-level objects (the laws of nature, the pictures) doesn't reduce straightforwardly to similarity at the lower level. The difference is that we have a grasp on similarity of pictures, but we don't have a grasp on what similarity of the laws is.

The only thing we *can* say about l-similarity is that it had *better not* straightforwardly supervene on the Humean mosaic of particular local matters of fact – it *better not* straightforwardly supervene on m-similarity. Why? If it did, the deviant world would come out as closer to the actual than the standard, and thus the counterfactual we were discussing would get the intuitively wrong truth value.

Lewis seems to suggest that, in reducing counterfactuals to l-similarity of possible worlds and in reducing laws to the mosaic, he has offered a full account of the truth values of counterfactuals.

> I take [counterfactuals] to be governed by similarity of worlds. . . To the extent that this similarity consists of perfect match in matters of particular fact, it supervenes easily on the arrangement of qualities; and to the extent that it consists of (perfect or imperfect) conformity by one world to the laws of the other, it supervenes if the laws do. (Lewis, 1987, p. xii)

However, despite these moves, we are no closer to understanding the right notion of similarity. Building on the analogy with the Hamming distance in the simple case, we can get a handle on m-similarity in terms of the number of properties that need to be modified at each point in the mosaic. And this kind of similarity seems to fit with Lewis' generally reductionist strategy as captured by his commitment to Humean Supervenience. However, as the preceding sections show, this is the wrong kind of similarity. And Lewis' analysis of lawhood does nothing to help understand what l-similarity is. All we know about l-similarity is that it doesn't supervene on m-similarity. It does however supervene on the "m-facts", so to speak. The general HS viewpoint seems to initially suggest that m-similarity would be the right similarity relation. And m-similarity has the advantage that we have some idea about what it would mean for worlds to differ more or less in terms of particular local matters of fact. It turns out that m-similarity isn't the right similarity relation for assessing the truth values of counterfactuals: Lewis suggests that l-similarity is. But we have much less grasp on what it means for worlds to differ more or less in terms of their laws of nature. And Lewis' reductive account of laws of nature does not help.

7 Exceptions and the big bad bug

We finished the last section by pointing out that we still don't really know anything about how similarity in terms of the laws of nature should behave. One thing that it would be difficult to argue with is that two worlds with *the same* laws should be more similar than worlds with different laws. Also uncontroversial is the suggestion that worlds with the same laws should be assessed for similarity by m-similarity (recall Lewis' hierarchy of aspects of similarity).

In the above we have been rather charitable to Lewis and discussed the best possible case for his view. Let's consider two arguably necessary refinements to Lewis' picture of laws that make things *even harder* for his account of counterfactuals. First, consider the case of a *ceteris paribus* law: a law that has a clause that rules out unspecified confounding factors or exceptions to the law. Such a law could be a true generalisation (let's imagine there's a way of filling out this picture that doesn't make lawhood trivial). Now consider our example again. The event of the ball floating still and not falling when let go of might be exactly the sort of exception that the ceteris paribus clause of the law deals with. In such a situation, it is possible for the three possible worlds we have been discussing to have *the same* laws. Now if all three worlds have the same laws, then there's *no way* for Lewis to make the intuitively true counterfactual come out as true. If anything, it will be false, since the deviant world is as l-similar as the standard world, and strictly m-closer than it.

Second, consider chancy laws. Again, chancy laws allow exceptions, and thus our example worlds might all have the same laws and the same conclusion follows. In either case, no amount of playing with the similarity relation is going allow us to recover the intuitively correct truth value of the counterfactual. David Lewis was a systematic philosopher, and his system was predicated on the idea of Humean supervenience. HS's Achilles heel was the notion of chance; Lewis called chance the "big bad bug". He worried that incorporating chances into his system could well pull the whole edifice down.

> There is one big bad bug: chance. It is here, and here alone, that I fear defeat. But if I am beaten here, then the entire campaign goes kaput. (Lewis, 1987, p. xiv)

His worry appears to be well-founded.

Note how these examples are working: laws of nature now admit of exceptions (for whatever reason). We think of the deviant world as a world with the same laws where the happenings discussed in the consequent of the counterfactual are an exception to the law. Such a world *has to be* closer to the actual world than the standard world, since actual and deviant share the same laws, and even if standard also shares the same laws, it is m-further away from the actual than the deviant world is.

Hawthorne (2005) points to something like this problem, and Williams (2008) provides a solution that involves modifying the similarity relation to allow worlds with the same laws to be dis-similar enough to block the deviant worlds from being salient.

The last section concluded that a robust analysis of similarity with respect to the laws was missing from Lewis' program. This section suggests that l-similarity in fact needs to be replaced with something more sophisticated.

8 Conclusion

Lewis' account of counterfactuals involves appeal to similarity between possible worlds. The original – and arguably intuitive – measure of similarity gives the intuitively wrong truth-values to counterfactuals. So Lewis offered an alternative account of similarity that appealed to laws of nature, and to the size of the "violation" of the laws required. Lewis argued that the violations of the laws required for standard possible worlds are smaller than the violations required for the deviant possible worlds. To return to our ball example, the violation that you let go of the ball rather than keeping hold of it is smaller than the violation that you let go of the ball and it just floats there. The intuition is that in the former case, it's just a tiny difference in neuron firings (or something like that) that needs to be different, and then everything else follows the laws, e.g. gravity. Contrast this with the deviant world where you let go of the ball but it floats where it is. This involves a bigger and more sustained violation of a law of nature.

This is an intuitively appealing picture of the situation, but it is not friendly to the view of laws that Lewis (and many others) endorse according to which laws are (some sort of) summaries of particular matters of local fact. Specifically, the intuitively appealing picture of the situation above clearly makes tacit recourse to the idea that laws *bring about* that the ball falls when let go of. But some such 'bringing about' is what Lewis' account of laws denies: the laws of a possible world don't *bring about* that the ball falls when let go of; the laws merely summarise those particular things that happen. So under a Lewisian account of laws it is not the case that in the deviant world involves "more" violation along the lines that the ball should be subject to the law of gravity. Rather the ball stays where it is and the laws of this deviant world had better reflect that fact. Given this, we don't see much reason to think that a (HS) laws-based hierarchy of similarity is going to make the standard possible world come out as closer than the deviant one.

To summarise, Lewis needs to appeal to the notion of l-similarity – similarity in terms of the laws of the possible worlds – in order to ground the correct truth conditions for counterfactual conditionals. Such a notion of similarity doesn't supervene on similarity at the level of the Humean mosaic – m-similiarity – and thus is unfriendly to Lewis' general project of a metaphysics based on Humean Supervenience. His two accounts are incongruent. The notion of l-similarity is underspecified, but one thing that we do know about it is that it *cannot* supervene on m-similarity on pain of getting the truth conditions wrong for counterfactuals. Further, if laws admit of exceptions – because of chances, or because of ceteris paribus clauses – then things are much worse: it becomes difficult to see how to salvage the similarity account.

How to avoid the problem? 0) Bite the bullet and deny our intuitive judgments about the truth-values of counterfactuals. 1) Give up the account of counterfactuals in terms of the relative similarity of possible worlds to the actual world. 2) Give up the account of laws of nature (either entirely or by dropping HS). 3) Further specify the similarity measure.

Let us make some remarks about these potential solutions. 0 seems to us to be too radical a route. No doubt not all of our intuitions are to be upheld, but we assume without further argument that (i) *is* true. What about 1? Lewis' account of counterfactuals is well-established and widely used. Therefore giving this account up would only be acceptable if there were compelling alternative accounts of counterfactuals, but to our mind there aren't. So 1 does not look much better. One could give up on the Lewisian account of laws and have a modally thick "oomphy" picture of lawhood. Of course this would be very un-Lewisian but given the many alternatives to BSA, 2 at least seems more promising than 0 or 1. However, it seems to us that 3 is the most promising way to avoid the problem given commitement to the truth of the intuitively true counterfactual and given the individiual appeal of Lewis' accounts of counterfactuals and laws. But just how to modify the similiarity measure whilst doing justice to both remains an open problem.

References

Bennett, J. (1984). Counterfactuals and temporal direction. *Philosophical Review 93*, 57–91.

Cohen, J. and C. Callender (2009). A better best system account of lawhood. *Philosophical Studies 145*, 1–34.

Fine, K. (1975). Critical notice of *Counterfactuals*. *Mind 84*, 451–458.

Hawthorne, J. (2005). Chance and counterfactuals. *Philosophy and Phenomenological Research 70*, 496–405.

Lewis, D. (1973). *Counterfactuals*. Blackwell.

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

Lewis, D. (1987). *Philosophical Papers Volume II*. Oxford University Press.

Lewis, D. (1994). Humean supervenience debugged. *Mind 103*, 473–490.

Stalnaker, R. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in Logical Theory*. Blackwell.

Williams, J. R. G. (2008). Chances, counterfactuals, and similarity. *Philosophy and Phenomenological Research 78*, 385–420.