

## POSSIBLE WORLDS OF DOUBT

### Abstract

A prominent contemporary anti-skeptical strategy, most famously articulated by Keith DeRose, aims to cage the skeptic's doubts by contextualizing subjunctive conditional accounts of knowledge through a conversational Rule of Sensitivity. This paper argues that this strategy courts charges of circularity by virtue of its selective invocation of heavy counterfactual machinery. Because of the danger that this invocation essentially employs a metric for modal comparison that is implicitly informed by judgments of epistemic sameness, this metric proves objectively indefensible. We have reason to fear that this metric is selectively cherry-picked in advance to support the very anti-skeptical conclusion for which the contextualist longs.

### I. Introduction

Because this paper concerns both “external world” skepticism and epistemic contextualism, it is useful to briefly describe each position. In the case of the former, this is relatively easy. Skepticism for present purposes, is the position that worldly knowledge is unattainable because various skeptical scenarios (involving *malign genies* and the like) cannot be effectively countered. In the case of contextualism, however, a brief description requires much more caution. The oft-cited reason is branding. Contextualism now comes in so many fragrances and flavors that we must step very lightly in any attempt to characterize a generic contextualist position. One chief branding distinction is that between *attributor* and *subject* contextualism, as we decide whose context is to determine the fortunes of knowledge claims (DeRose 1999, 190-191). Another is that between *conversational* and *non-conversational* contextualism, as we decide which features of context properly count as determinative (Williams 2004, 193). A third is that between varying accounts of distance, as we decide the nature of the changes that are ushered in by distinct epistemic contexts. It is thus that we are confronted by “contextualists” as different as Keith DeRose and Michael Williams (DeRose 1995; Williams 2001). Although the former insists that it is the conversational circumstances of attributors of knowledge that render various possibilities relevant or irrelevant by raising or lowering standards of knowledge along a single scale of severity, the latter maintains that it is the speaker's

background information and practical interests that determine her standards of knowledge (and, in cases of radical skeptical challenge, whether or not a recognizably ordinary “knowledge relation” is even in play at all). In the light of such diversity it can be quite misleading to speak of “contextualism” as a generic position.

For this paper's purposes, however, this complication can be ignored. By “contextualism” we are targeting only the classic version described by Keith DeRose in "Solving the Skeptical Paradox" (as well as any others relevantly similar to it), a position we label Classic Conversational Contextualism (henceforth, *CCC*)<sup>i</sup> (DeRose 1995). My reason for this is a critical one: CCC is particularly susceptible to charges of circularity due to its distinctive invocation of heavy counterfactual machinery. This is because such invocation threatens to employ a metric for modal comparison that is implicitly informed by judgments of epistemic sameness. Section II details the contextualist's anti-skeptical project. Section III lays the groundwork for my own argument. Section IV develops the argument itself. In Section V we canvass a number of potential objections and replies.

## **II. Contextualizing Externalism to Anti-Skeptical Ends**

CCC may be usefully viewed as an effort to identify the proper conditions for applying the knowledge predicate by appealing to particularizing circumstances that extend far beyond those commonly invoked by classic Goldman-style externalism (Goldman 1967, 1976). For such externalists, knowledge typically accrues as a function of causal or reliabilist processes that systematically relate knowing agents to the states of affairs that render their beliefs true. From the standpoint of CCC, such unadorned externalism is inadequate for the task of fully regulating ascriptions of knowledge. There are two reasons for this. The first is that such externalism effectively assumes the invariance of context. A person knows or doesn't know, on such accounts, as a function of his placement within a network of causal relations (in particular, those causal relations that serve to render belief-forming processes reliable), irrespective of whatever individuating conversational context he may happen to occupy. The second reason is that such

unadorned externalism is much too narrowly naturalistic. Subjunctive conditional analysis is required for its inclusion of a modal element. DeRose's broadest constructive strategy is that of contextualizing an externalistic, but non-empiricist, base account. To this extent, we may usefully think of CCC as an "externalism-plus" position. Although CCC can hardly treat unadorned contextualist criteria as irrelevant to knowledge attribution, it must invariably deny that these criteria allow us the resources with which to tell a complete story about when such attribution is appropriate.

The non-empiricist base account on which DeRose builds is that of Nozick, for whom knowledge is true sensitive belief (Nozick 1981). On Nozick's penultimate telling (taking *M* to be a constant belief-forming method), *S* knows *P* if and only if *P* is true, *S* believes *P*, *S* would believe *P* (on the basis of *M*) if *P* were true and *S* wouldn't believe *P* (on the basis of *M*) if *P* were false (Nozick 1981, 179). DeRose aims to contextualize this subjunctive conditional account through a conversational Rule of Sensitivity, according to which the possibilities to which we hearken become relevant by virtue of said hearkening, and pertinent standards for knowledge inflate or deflate in response to our very articulations of our epistemic inquiries (DeRose 1995, 35-38). In this way, DeRose aims to build the sensitivity requirement, not into the concept of knowledge, but into a background account of how the standards for knowledge are conversationally raised and lowered (DeRose 1995, 38-39). More formally, DeRose writes, "When it is asserted that some subject *S* knows ... some proposition *P*, the standards for knowledge tend to be raised ... if need be; to such a level as to require *S*'s belief in that particular *P* to be sensitive for it to count as knowledge" (DeRose 1995, 36). Alternatively stating things in possible worlds' locution, DeRose writes, "When it's asserted that *S* knows . . . that *P*, then, if necessary, [one must] enlarge the sphere of epistemically relevant worlds [to be considered and ruled out] so that it at least includes the closest worlds in which *P* is false" (DeRose 1995, 37).

The resulting account, DeRose tells us, explains why skeptical claims of ignorance have the force they do in the conversational contexts within which they appear, while still failing to imply our ignorance of the ordinary knowledge claims to which we rightfully and stubbornly

adhere. Though closure (contra Nozick) and skeptical doubts persist, our ordinary epistemic convictions remain secure, because they are issued from a context in which fewer counterpossibilities need be raised and addressed before knowledge claims can be issued. We may not know that we are not a genie's plaything, but this bit of philosophical ignorance fails to render us ignorant in ordinary contexts. That is, it fails to imply the falsity of prosaic knowledge claims within any context from which we are ordinarily inclined to assert them.

Applied to skepticism, this account yields the following story. In ordinary contexts, as opposed to skeptical ones, we assess the claim to know  $p$  by advancing our perspective, accordion-like, to the nearest possible world in which  $p$  is false. From there we apply the very same belief-forming method that we employ in the actual world (e.g., "consultation with visual experience"). Consider Moore's claim to know that he has hands. In ordinary contexts, where skeptical scenarios are not entertained, we are to assess the counterfactual, "If Moore didn't have hands, he wouldn't believe that he did," by viewing things from the nearest possible world in which his hands are missing. This possible world, we are assured, is one in which Moore is handless by virtue of some perfectly ordinary possibility (e.g., his hands have been lost to some tragic industrial accident), rather than by virtue of some extraordinary possibility (e.g., involving *malign genies* or the like). But in skeptical contexts things are very different. In these contexts, DeRose's Rule of Sensitivity dictates that we consider things from the standpoint of the nearest possible world in which Descartes' belief that he is not being deceived by a genie is false. And in this world, given the genie's MO of rendering herself undetectable in principle, Descartes cannot be said to know that he is not being deceived. Moreover, in ordinary contexts the inference from "Moore knows that he has hands" to "Moore knows that he is not the (handless) victim of a *malign genie*" also goes through, just as long as only the former knowledge claim is at deliberative issue. This is because the nearest possible world in which Moore is deprived of hands is similar enough to the actual one to keep the genie from her post, as it were, thus preventing her from sabotaging Moore's knowledge of his own handedness. Let's call the type of anti-skeptical strategy at play here a *CCC anti-skeptical strategy*,

its proponent the *CCC anti-skeptical strategist*, and the argument resulting from its application the *CCC anti-skeptical argument*.

In summary, the above account aims to be one on which contextualism is grounded not only in thoroughly externalist criteria, but in externalist criteria of a sort that promise, by virtue of their modal richness, to render them distinct from the unadorned externalist criteria of classic Goldman-style reliabilism. On this account, skeptical claims (e.g., that someone does not know that he isn't some genie's benighted plaything) are rendered true by the distinctively high standards of "philosophical" knowledge, and ordinary claims (e.g., that a person knows he has hands) are rendered true by the looser standards of prosaic epistemological parlance. Ignorance is segregated to philosophical contexts of knowledge, and the rule of sensitivity explains why this is the case.

### **III. Critical Strategy and Framing Assumptions**

To repeat, this paper's broadest goal is to question the legitimacy of the modal appeals at play in CCC anti-skeptical strategies. Arguments motivating this doubt debut in Section IV. First, however, we need to engage in a little stage setting and say a few things by way of motivation, vernacular, strategy and qualification.

What we can't over-emphasize is that that the arguments in Section IV do not stem from any *general* distrust of either modal contrivances or their employment in other areas of philosophy (e.g., toward the end of analyzing agency, laws, probability or causation). Irrespective of whether or not counterfactual explications of philosophical notions are generally suspect, they are particularly suspect in the context of CCC anti-skeptical strategies. The reason for this concerns the status of the modal metric that must be in place whenever we endeavor, toward anti-skeptical ends, to zone neighborhoods of possible worlds into those nearer and farther out. Such modal metrics are prone to objective indefensibility because they are likely to be selectively cherry-picked in advance to support particular anti-skeptical agendas.

To motivate the objection at issue here, let's ask a blunt question. Why should we believe that the nearest possible world in which Moore doesn't have hands is one in which they have been

removed by an industrial accident, say, rather than one in which he is the handless unknowing player in some radical skeptical scenario?<sup>ii</sup> What objective criteria of modal similarity could render the former “way things might have been” objectively closer than the latter “way things might have been” to the actual world? Any given possible world is similar to the actual world in some respects and different from it in others. And even if it proved possible to compare raw numbers of similarities and differences, no such numerical comparisons could themselves ever be dispositively conclusive. To determine our relative proximity to various possible worlds in a manner congenial to CCC anti-skeptical strategies, we would need to invoke standards that are up to the task of objectively dictating which respects of similarity and difference matter more than others, where such objectivity requires that these standards be defensible without appeal to any preferential biases we might have toward worlds in which our ordinary knowledge-acquisition procedures prevail. The reason for this is that CCC anti-skeptical strategies are fueled by the intuition that it is the radical skeptical scenarios that define those worlds most metaphysically distant from our own. Thus, we certainly cannot assume from the outset that it is these scenarios’ knowledge-inhibiting status that renders them metaphysically remote.

For ease of expression, let’s use the following vernacular. Where *metaphysical similarity* or *distance* registers the degree of difference *simpliciter* between various possible worlds, *epistemic similarity* or *distance* registers the degree to which each world’s standard epistemic claim-making procedures fail to be veridical or knowledge-acquiring in the other. In these terms we can now state our challenge more succinctly. *To cogently employ CCC anti-skeptical strategies, we must show that the likes of “industrial accident” worlds are, as a kind, metaphysically closer to actuality than are the likes of “skeptical scenario worlds,” and that said metaphysical distance tracks epistemic distance without being implicitly calibrated by it.* For ease of expression, let us refer to this claim (that such objective measurement and independent tracking occur to anti-skeptical ends) as the *Contextualist Objectivity and Independence of Metaphysical Distance Metrics Thesis*, or *COIM* for short. This acronym, though ugly, is at least pronounceable. The concern of this paper, then, is to

argue against the plausibility of COIM, or at least against those who would accept COIM as the default position without supporting justification.

As a prelude to Section IV, we need to address three additional points before turning to actual arguments. The first concerns the presuppositional baggage that informs our discussion of similarity. The second concerns the dialectical position that our argument occupies within the debate it seeks to motivate. The third concerns methodology.

As for presuppositional baggage, we would do well to pursue a policy of presumptive minimalism when dealing with questions of similarity and difference *in general*, and with questions of similarity and difference between possible worlds *in particular*. With regard to similarity *per se*, we would do well to understand similarity as a function of identity. Two objects are similar to the extent that they are qualitatively identical in various constitutive respects. We would then do well to remind ourselves of the philosophical chestnut that, although debates about identity may indeed be intelligible, this intelligibility does not typically derive from concerns about identity *simpliciter*, but rather from questions about the idiosyncratic identity conditions of the distinct and sundry kinds at issue. Thus, to ask about personal identity is to ask little, if anything, about the concept of identity itself, but *everything* about the concept of personhood. With this in mind, we may now note a simple but useful truth. Questions that arise concerning identities between domains substantially depend upon our characterizations of these domains. Are there composite objects, or merely the basic parts thereof, about whose identity conditions we must be concerned? About how many levels of things must we ask whether they endure or perdure through time? To whatever extent allowable, we should ignore these and related questions as tangential to our main concerns. Perhaps we can do this to some extent by taking our substantive world-inhabiting items of comparison to be, not natural objects and artifacts, but points in space-time and regions thereof, to and between which properties and relations may be, respectively, ascribed. For present purposes, let's provisionally concede that such points, regions and their time-stamped exemplifications of properties and relations (i.e., facts) exist, whatever else may exist, and in this way set aside as many questions as

we can concerning mereology, emergence, identity over time and the numbers and classes of threatening counterparts.

With regard to similarity and difference between possible worlds, more specifically, we would do well to briefly position our concerns in relation to the literature. This tradition traces to Stalnaker and Lewis, both of whom investigate possible worlds' similarity in the course of examining counterfactual conditionals as tools with which we might logically characterize the features of actuality. (Stalnaker 1970; Lewis 1973a, 1973b). For both authors, conditionals of the form, "If A were true, then B would be true" apply in the actual world only when B accompanies A in most similar non-actual worlds in which A applies. That is to say, such conditionals are true when less departure from reality is required to make their antecedents true in conjunction with their consequents than is required to render their antecedents true alone. Beyond this, details vary, as Stalnaker and Lewis offer competing formalisms with which to characterize those worlds differing minimally from our own. Stalnaker targets this class through a selection function designed to pick out that world in which A obtains and approximates actuality as closely as A's truth allows (Stalnaker 1970). Lewis proposes an ordering metric on which any two non-actual worlds must either be equidistant from actuality or else line up in progressive serial order with respect to it. He then stipulates that "if A were true, then B would be true" applies if and only if either no worlds accommodate A, or else some world accommodating both A and B is closer to actuality than is any world accommodating A but not B (Lewis 1973a, 1973b, 1986). Various differences follow from these differences, in turn. Lewis jettisons both Stalnaker's uniqueness and limit assumptions, which collectively dictate that exactly one world is metaphysically closest to actuality. Thus, Lewis' worlds may tie for similarity, but not Stalnaker's. For Stalnaker, each world has one and only one most similar neighbor.

Beneath these differences, however, the thematic similarities between Lewis' and Stalnaker's accounts prevail, at least insofar as our present concerns go. In particular, both authors effectively treat possible worlds' comparative similarity as an intuitively familiar one, and so provide little method for judging inter-world distance other than simple reflection upon the relevant

counterfactual conditionals themselves. In Lewis' mind, such reflection leads to particularly pregnant suggestions, as he notes ways in which similarity relations between worlds may be very different from similarity relations between objects inhabiting worlds. For, if we want to say that "If Nixon had pushed the button, there would have been a nuclear holocaust" is true by virtue of the latter event's occurrence in most similar non-actual worlds accommodating the former event, then we must find some grounds for ignoring the overwhelmingly obvious fact that post-apocalyptic scenarios certainly seem to be metaphysically very distant, indeed, from our actual non-radioactive environs.<sup>iii</sup> Lewis' response to this problem is to suggest that the similarity relation invoked by counterfactuals requires that those worlds most similar to our own agree with actuality only up until shortly before the antecedent occurs, whereupon a "localized" and "inconspicuous" miracle occurs to accommodate the antecedent, from which point on the world progresses in full compliance with actual law (Lewis 1979, 472).

Thus, when comparing possible worlds, Lewis urges, we must carefully note the tradeoffs we make between differences in world-governing laws and differences in world-constitutive facts. Toward this end he proposes ranked desiderata, the effect of which is to promote, in lexical order, *first*, widespread agreements in regulative natural laws; *second*, widespread agreements in constitutive facts across extended spatiotemporal regions; *third*, agreements in localized regulative natural laws; and *fourth*, agreement in particular constitutive facts, "even in matters that concern us greatly" (Lewis 1979, 472). Put less carefully but more succinctly, Lewis maintains that big miracles are more difference-making than little miracles, which are less difference-making than whole-scale variations in constitutive facts across broad contiguous swaths of reality. In his analysis, the worlds most similar to actuality are those in which actual laws of nature are never violated. However, precise similarity of fact across significant spatiotemporal regions can outweigh such nomological similarity if it can be achieved at the cost of only small, localized miracles.

Lewis' lexically ranked list of metaphysical similarity desiderata remains controversial, and we invoke it here mainly for the cautionary moral it provides: confidence is misplaced when we endeavor to provide such lists.<sup>iv</sup> Even so, we would do well to keep in mind the major question

that prompts them. When assessing similarity and difference between possible worlds, to what degree should we emphasize the preservation of law over fact? Much depends on how we answer this question and, even more pointedly, whether or not we even *can* non-arbitrarily answer it one way rather than another.

The second point we need to address before turning to actual argument concerns the dialectical position that said argument occupies within the debate that it seeks to motivate. It is important, within the context of this debate, to remember that it is *we* who have the easier position to defend. The reason is that it is the CCC anti-skeptical strategists, not us, who are committed to making a *very* strong pair of assumptions. These are the assumptions that (1) metaphysical distance is amenable to non-arbitrary measurement, and (2) said metaphysical distance tracks epistemic distance without being implicitly calibrated by it. As with any assertion that specified similarity spacing exists *with* a specified critical significance, it is incumbent upon COIM's proponents to defend their position against the charge of wishful thinking. Think of the history of racial classification, throughout which the perennial impulse has been to justify prejudices by invoking distinctions clouded by those very prejudices from the outset (Gould 1981). The burden of proof is on the CCC anti-skeptical strategists to counter this charge of wishful thinking. Our goal here needs to be only to provide reasons to doubt that any such countering defense can be provided in *systematic* and *principled* terms.

The third point that merits mention before we proceed to actual arguments concerns the level of generality at which we should reason about these matters. In particular, we should seek to avoid, to whatever extent possible, deriving overly general morals from the consideration of particular and idiosyncratic cases. The dangers of such a procedure are hinted at by Stalnaker (as told to Bennett), who writes of how, in all such cases, “intuitions are not decisive” because of how “theoretical considerations mix with intuitive judgments” (Bennett 2003, 193). Even Lewis, despite his thoroughly case-driven methodology in these matters, recognizes the problem when he writes of how “our respect [for] the extreme shiftiness and context-dependence of similarity” should make us very wary of making “offhand similarity judgments” and then “assuming that ... they will do for all

purposes” (Lewis, 1979, 466). Because the suspicion we aim to motivate in this paper is that the modal metrics presupposed by CCC anti-skeptical arguments are selectively cherry-picked in advance to support anti-skeptical agendas, it is essential to the cogency of such arguments that the counterfactual similarity criteria they employ be reasonably objective, rather than being derived from the specific programmatic aims and purposes that motivate them. Once again, we should put the onus on CCC anti-skeptical strategists to defend their modal metrics in *systematic* and *principled* terms. A handful of cherry-picked thought experiments will not do.

#### **IV. The Argument from Modal Circularity**

With the preceding locutions, assumptions, clarifications, qualifications and disclaimers in place, let us pose our blunt question once again. If we take our constant method of belief formation to be something like “the method of consulting sensory experience,” why should we suppose that the nearest possible world in which Moore doesn’t have hands is one in which they have been removed by an industrial accident (an “*IA world*”), rather than one in which he is a handless occupant of some radical skeptical scenario? (To render our comparisons concrete and telling, let’s henceforth restrict “radical skeptical scenarios” to “brain in vat” scenarios, thus taking our *malign genie*, occupying *BIV worlds*, to be the non-magical mad scientist of recent philosophical vintage, maliciously poking at his victim’s brain brain, after having envatted it mere moments before.) Is it obvious that every IA world is closer to actuality than every BIV world in both its laws and its attributions of properties to and relations between constitutive space-time regions? At first blush, it may seem that a slam dunk argument for this conclusion is surely waiting in the wings. Certainly, when we consider alternative worlds at the very moment when our knowledge of our handedness becomes at issue, BIV worlds appear to be much more broadly different from actuality than are IA worlds.

However, this claim may be too rash. To see why, we need to canvass the worlds we are comparing in the broadest possible terms. Like Lewis, we can do this by considering two general possibilities. With the first possibility, the past histories’ of the worlds we consider stand in

*indeterministic* relation to these worlds' present conditions, either wholly or in part. With the second possibility, the past histories of the worlds we consider stand in *deterministic* relation to these worlds' present conditions, and thus become realms from which Lewis' miracles, both large and small, are excluded. Note, however, that "determinism," as understood here, is not quite Lewis' commonplace notion, because it requires merely law-governed ordering *rather than* unique predictability. Natural law in a such a "deterministic" world may constrain outcomes in a merely probabilistic manner.<sup>v</sup> To be deterministic, for our purposes, a world be merely law-governed.

Consider the first possibility. Because indeterministic worlds are unconstrained by any need to accommodate causally compatible background stories, they offer us limiting scenarios which maximize IA and BIV worlds' similarities to actuality, because both the laws regulating them and the factual conditions constituting them remain identical to actuality and to each other until the last possible moment. In such indeterministic settings, *every* point in the inclusive history of an IA world up to the moment of amputation, and *every* point in the inclusive history of a BIV world up to the moment of the genie's deception, may be identical to those of the actual world because, in the absence of regulative laws, past constitutive facts are not constrained by nomological prohibitions on what may follow them in the future.

When comparing such indeterministic cases, we should realize that our strongest intuitions are that IA worlds, as a class, enjoy greater metaphysical proximity to actuality than do BIV worlds. This suspicion is well-founded, for only in these instances (in which a radical determinative disconnection separates past and present events) can we ground our intuitions concerning similarity and difference upon single snapshot observations of reality at particular moments in time. There is a telling reason that this is the case. In the absence of tracking laws and histories, only last-minute mass-reshufflings of constitutive facts cause BIV worlds to diverge from IA worlds, and from actuality itself. This boosts our trust in COIM because the last-minute reshufflings required to turn IA worlds into actuality certainly appear less extreme than those required to turn BIV worlds into actuality. Industrial accidents occur all the time, after all; but BIV worlds are the stuff of philosophy and science fiction. My being an inattentive machinist requires merely that my

cognitive habits and profession vary. My being an envatted brain requires the truth of numerous contrary-to-fact propositions concerning both technical advances and the actions and motives of sadistic implementers.

Now consider the second possibility, to which determinism applies, and which excludes both large and small miracles. What bears emphasis is that the assorted pertinent differences in these cases between various alternative worlds and our own cannot date merely from the single selective moment in time when a person's knowledge of his handedness becomes an issue (or the moment when the radial arm saw slices off his arm, or when the mad scientist's machinations are put into motion). That is, when comparing deterministic worlds, we are not allowed to simply choose or stipulate the moment at which similarities and differences begin to accrue. The fact that every moment in a deterministic world is embedded within an inclusive cosmic history must invariably imply that other differences have been occurring throughout the causal span of that history. And these are differences that we are obliged to tally and ponder in the course of examining and considering contrasting worlds. Thus, the observation that BIV worlds differ from actuality more than IA worlds do at the moment when knowledge becomes an issue, *in itself*, tells us nothing. For it to imply our own greater metaphysical distance from BIV worlds, it must work in conjunction with the presupposition that smaller overall differences in the present invariably require smaller causally originative differences in the past. *But how can this be assumed?* If we concede that numerous different networks of causal antecedents may conspire to bring about a given single state of affairs, then we have no *systematic* reason to suppose that the causal antecedents of occurrent IA worlds differ less from actuality, overall, than do those of occurrent BIV worlds.

Let us now assume two things. First, let us assume that the actual world itself is (uniquely or probabilistically) deterministic, governed by whatever regulative principles actually hold sway. Second, let us assume the first member of Lewis' lexically ranked list of modal similarity desiderata, according to which cross-world similarity requires widespread agreement in natural law. From these two assumptions it follows that the worlds metaphysically nearest to actuality are also

deterministic, and therefore subject to principles of law-like ordering much like those that characterize reality. Both of these assumptions are reasonable. The first is reasonable because “determinism,” on our telling, accommodates both classical and quantum models of reality. The second is reasonable because widespread regulative nomological divergence is unlikely to be offset by widespread constitutive factual agreement, because the former is prone to carry so much factual variation in its wake. That is, differences in law are very likely to sponsor their *own* differences in subsequent fact, making worlds that nomologically vary from actuality different from it along two dimensions, rather than one.

Taken successively, these two claims are significant; the first for its ability to render cross-world similarity comparisons manageable, and the second for its ability to cast doubt on COIM. For if deterministic worlds are, as a rule, more similar to actuality than are indeterministic worlds, then we need to consider only the former when scanning the modal horizon for our nearest metaphysical neighbors. And if, as argued above, IA worlds impress us as more clearly similar than BIV worlds to actuality only when we compare indeterministic possibilities, then we may yet hope to circumvent our initial sense of assurance that COIM is a slam dunk position. Considered together, these two assumptions provide both focus and motivation to our inquiry. When asking if COIM is truly unproblematic, we are now allowed to restrict our attention to the following question. Considering *only* those deterministic possible worlds governed by laws broadly similar to our own (in effect, allowing only for Lewis’ “minor miracles”), is it clear that the metaphysically nearest such world in which Moore doesn’t have hands is one in which they have been removed by an industrial accident, say, rather than one in which he is some poor pitiful and envatted science fair project?

To some, perhaps even this more restricted question may seem to have an immediate answer. That is, it may still seem obvious to some that BIV worlds require much more variation from actuality in constitutive fact and/or regulative law than do IA worlds, allowing COIM to remain, slam dunk position or not, at least very likely true. The reasoning underlying this conviction is easy to discern. BIV worlds are worlds in which the causal/perceptual relations

informing intra-world reliabilist relations must not only fail, but must fail *spectacularly*. For these are worlds in which people's beliefs about their environs are brought about, not by the actual mechanisms we take to support them (whereby tables' reflectance features cause table percepts), but by bizarre alternative mechanisms (whereby direct neural stimulations cause table percepts). Perhaps this is enough in itself to render probable the contention that epistemic distance closely tracks metaphysical distance, and thus the contention that BIV worlds are invariably more distant from actuality than are IA worlds.

However, broad theoretical grounds still caution us to remain wary of this conviction. *The knowledge relation is merely one among many that collectively constitute our position within the world we occupy.* It is for this reason that doubts should remain. For why should we suppose that the specific worlds in which our belief-forming mechanisms remain reliable are the ones that lie metaphysically closest to our own? *Doesn't this place epistemic similarity between possible worlds above all the other possible respects in which such worlds may be the same?* Absent some reason to think that epistemic similarity is especially significant, we would seem to be barred from privileging it *vis-à-vis* the many other possible ways in which worlds may be metaphysically alike. But what, if anything, could produce such special significance? Put differently, what could allow us to assume COIM?

Let us call cross-world variations in properties and relations that directly figure into the reliability of beliefs *epistemic differences*, and those that do not so figure *non-epistemic differences*. Correlatively, let us call properties and relations that directly factor into epistemic differences *epistemic difference-making*, and properties and relations that do not so factor *non-epistemic difference-making*. There seem to be three general ways that COIM might prove ineluctable. First, it might turn out that epistemic difference is naturally *more encompassing* than other kinds of difference, directly constituting more in the way of minor miracles and differences in constitutive fact. Second, it might turn out that epistemic difference *causes or determines* more in the way of concomitant non-epistemic variation than does non-epistemic difference. The idea here is that the sorts of properties and relations that render our belief-forming mechanisms

reliable might play a special causal or determinative lynchpin function, allowing them to result from or bring about particularly numerous and diverse differences in other properties and relations, to which they may initially seem only marginally, if at all, connected. Finally, it could turn out that epistemic difference is just *intrinsically more significant* than other kinds of difference, and thus counts for more than other kinds of difference in determinations of metaphysical proximity. Let us briefly examine each of these possibilities in turn.

The first possibility, again, is that epistemic difference might be intrinsically more encompassing than non-epistemic difference, by *constituting* more in the way of minor miracles and differences in constitutive fact. But it is hard to see how a principled defense of this position could be made. One reason is that there is clearly at least one respect in which BIV worlds must *always* enjoy greater metaphysical proximity to actuality than IA worlds, namely, in their narrative details concerning the phenomenological and cognitive lives of the agents we imagine being deceived. BIV worlds are set up to create illusions of normalcy, after all. Within them, Moore's *experienced* life is devoid of amputation, and his subsequent experiences and thoughts run parallel to those of actuality in consequence.

A second reason to eschew the assumption that epistemic difference is more intrinsically encompassing than non-epistemic difference is that BIV worlds, when invoked to the ends of CCC strategic deliberation, are arguably *less* likely than IA worlds to involve minor miracles, at least when we accept the likes of Lewis' lexically ranked list of metaphysical similarity desiderata. Why is this? Remember again how Lewis aims to employ his ranking criteria. As he envisions things, the worlds most metaphysically proximate to actuality are those that agree with it fact-wise until just shortly before the pertinent conditionals' antecedents apply, at which point "localized" and "inconspicuous" miracles occur to accommodate these conditionals' antecedents, and from which point on the world progresses in full compliance with actual law. Thus, if we accept Lewis' priority governing modal similarity ranking *and* the CCC strategists' assessment of BIV worlds as further from actuality than IA worlds, we get some rather interesting results. In particular, we get the result that IA worlds, for being closest to actuality, are settings where pertinent antecedent events (e.g.,

losing hands to a circular saw) are likely to occur, not by virtue of variant initial conditions, but by virtue of passing fluctuations in local determinative law. BIV worlds, on the other hand, are allowed to remain realms where deceivers come to do their mischief by virtue of variant initial conditions alone. BIV worlds are non-magical, after all. Thus, to imagine them coming to pass we need only imagine them varying from reality in their constitutive determining historical facts.

Why should this fact bother us? After all, it only fuels suspicion that metaphysical and epistemic similarity criteria come apart *if* we already doubt that “minor” variations in regulative law are less difference-making than “major” variations in constitutive fact. Thus, we need to ask if this suspicion is warranted. I suggest that it is. The conviction that “minor” variations in regulative law are less difference-making, *simpliciter*, than “major” variations in constitutive fact stems, not from the first of Lewis’ ranked lexical imperatives (that proximate worlds must, first and foremost, enjoy widespread agreement in natural law), but from his second two desiderata (jointly dictating that factual agreement across extended spatiotemporal regions constitutes greater difference than regulative agreement in localized natural law). These second two desiderata, however, are clearly more questionable than the first, so we should be much more wary of them.

The main reason for this additional wariness is the following: to defend the second two ranked imperatives, we must compare apples and oranges in a way that we need not do when defending the first. Our defense of the first desideratum, remember, is that widespread regulative nomological divergence is unlikely to be offset by widespread constitutive factual agreement, because such nomological difference is liable *in itself* to carry much factual variation in its wake. Thus, our justification for the first desideratum does not require that we compare potential incommensurables. We need not decide that the preservation of laws, *per se*, is more important to metaphysical similarity than is the preservation of facts, *per se*. We need only note that widespread divergence of law entails widespread divergence of fact, but *not vice versa*. Lewis’ second two lexically ranked desiderata, on the other hand, require us to directly compare *major* changes in varying fact to *minor* changes in varying law, and to judge the former to be more

significant than the latter. But it is simply not clear how we are to objectively compare such different sorts of items, or how we are to contrast and compare the *degree* of a change in law and the *degree* of a change in fact along a single invariant scale of measurement. This is a problem that we might resolve through reduction. For instance, if we took laws to supervene on dispositions and dispositions to possess their causal powers essentially, then changes in laws might be seen as of a kind with changes in fact (Bird 2005). But to assume such tenets is to front-load our position with some rather substantive, and possibly unwelcome, background commitments from the start.

The second possible way (listed five paragraphs back) that COIM might prove unavoidable is that epistemic difference, though not *constituting* more in the way of variation *per se* than non-epistemic difference, might *cause* or *determine* more such variation. The former might be characterized by distinctive causal fecundity. The idea here is that the sorts of properties and relations that render our belief-forming mechanisms reliable might play a special causal or determinative lynchpin function, causing them to typically result from or bring about especially numerous and diverse differences in other properties and relations, to which they may initially seem only marginally, if at all, connected. This is the notion that a person's epistemic difference-making properties and relations to his surroundings may exercise a distinctive domino effect, changing or resulting from other world constitutive facts much more readily and completely than do differences in properties and relations that aren't epistemic difference-making.

However, it is hard to see how a principled defense of this position could be made, either. Again, remember that Lewis would have us believe that the IA world closest to our own is one where the divergence from actuality that robs Moore of his hands is a fleeting change in natural law, rather than a difference in causally efficacious preceding facts. Much of the above-described strategy's cogency for defending COIM depends upon whether or not we accept these ordering priorities. If we do, then a distinctive causal fecundity on the part of epistemic difference-making properties and relations might very well support IA worlds' greater metaphysical proximity to our own relative to BIV worlds, because the latter exhibit greater difference in precipitating initial

conditions from actuality. But, if we do not accept it, then no such support is needed, because we will have effectively denied that minor differences in law are less difference-making than major differences in fact. The point, once again, is that our candidate strategy for defending COIM is effective only if we can appropriately compare variations in law (apples) to variations in determining fact (oranges). But this requires us to invoke an objective metaphysical similarity metric which we can guarantee, in advance, will not diverge from the epistemic similarity metric that informs our anti-skeptical agenda.

The third possible way (cited seven paragraphs back) that COIM might prove unavoidable is that epistemic difference might just be *intrinsically more significant* than other sorts of difference, and thus count for more in determinations of metaphysical proximity. But, when asking how such a judgment could ever be made, we would do well to remember Lewis' own warning that similarity metrics are invariably subject to an extreme context-dependence that is directly determined by the purposes to which they are employed (Lewis 1986a, 466; 1973a, 67). There is something strongly suspicious about talk of one similarity metric being more or less *intrinsically* significant than another, particularly when such talk arises in the course of an effort to privilege a given similarity metric to a specific programmatic end. The temptation is very strong in such contexts to privilege precisely that metric required to support the moral for which we pine.

I suggest that this suspicion is particularly appropriate to our dealings with CCC-strategic reasoning. The threat of such circularity lurks constantly in the background. Consider, for example, the following excerpt from "Solving the Skeptical Paradox:"

Context, I've said, determines how strong an epistemic position one must be in to count as knowing. Picture this requirement as a contextually determined sphere of possible worlds, centered on the actual world, within which a subject's belief as to whether P is true must match the fact of the matter in order for the subject to count as knowing. (Given [previous] results . . . , we must again remember either to restrict our attention solely to those worlds in which the subject uses the same method of belief formation she uses in the actual world, or to weigh similarity with respect to the subject's method very heavily in determining the closeness of possible worlds to the actual world.) Call this sphere the sphere of epistemically relevant worlds. As the standards for knowledge go up, the sphere of epistemically relevant worlds becomes larger. The truth-tracking of one's belief must extend further from actuality for one to count as knowing. Given this picture, the Rule of Sensitivity can be formulated as follows: When it's asserted that S knows (or doesn't know)

that P, then, if necessary, enlarge the sphere of epistemically relevant worlds so that it at least includes the closest worlds in which P is false. (DeRose 1995, 20)

This passage provides the final setup for DeRose's first detailed statement of his "powerful solution" to the skeptical puzzle. Unpacked, the idea is as follows. "Epistemically relevant" worlds are those where knowledge requires both true belief and belief-forming methods identical to those of the actual world. So understood, epistemic relevance is used to group possible worlds for consideration in one of two ways. Either epistemically relevant worlds are subject to ordering by the CCC strategists' modal metric, *or* epistemic relevance is itself regarded as largely *constitutive* of that metric. Let us focus on the second option, because it is especially pregnant with the potential for abuse.

At first glance, it may seem innocuous to suggest that we treat as relevantly similar only those worlds across which a constant belief-forming method is employed. This may seem to be a harmless qualification required to address Nozick's "grandmother objection" and its subsequent requirement that beliefs track truth via the same method across variant possible worlds. (The "grandmother example" presents us with a case in which granny knows her grandson is well upon seeing him *despite* the fact that she would also feel confident about his health via an *alternative* epistemic method in a nearby possible world where he is, in fact, stone-cold dead.) But on deeper consideration, the proposal here should give us serious pause. The reason lies in epistemic relevance's radical context dependency, on DeRose's own telling.

When one searches for the possible worlds most similar to the actual world in which the grandson is not well, the respects in which the possible worlds are to resemble the actual world is a highly context-sensitive matter. Especially where the context focuses one's attention on the grandmother and her cognitive and recognitional abilities, one can place heavy weight upon similarity with respect to the method she is using to arrive at her belief, and then it can seem that in the closest world in which the grandson is not well, she's looking right at him and seeing that he's not well, and so does not believe he is well. (DeRose 1995, 21)

From the standpoint of our concerns in this paper, the use of "epistemic relevance" criteria to determine a modal metric is suspect whenever said use informs the similarity rankings of those very possible worlds wherein a putative knower's ability to track truth is at issue. *In the passage*

*above such use clearly threatens to do exactly this.* Why? Because DeRose's implicit suggestion is that we resolve the so-called "generality problem" (the problem of deciding how to specifically characterize a belief forming process) in terms that prejudice our possible world rankings from the outset. Suppose we characterize the grandmother's method of belief formation in the way casually suggested above, as one in which she physically inspects her grandson. Such a description of belief forming method dismisses BIV worlds from consideration from the start (given the eyeless condition of granny's envatted brain); whereas more minimalist descriptions of belief forming method (e.g., consultation with and assessment of "visual" experience) do not. The point is that the very sense of salience that CCC strategists invoke in the course of deciding which worlds are metaphysically closest to actuality by virtue of being most "epistemically relevant" is itself likely to be prejudiced by an implicit metric of epistemic similarity. And this, we have seen, is precisely what anyone seeking to defend COIM must avoid.

## **V. Objections and Replies**

However, a number of potential objections arise. In particular, might our reservations concerning CCC strategists' invocation of modal machinery seem both *too selective* and *too late*? To worry that such invocation is *too selective* is to worry that any reasonably comprehensive philosophical worldview is likely to require modal metrics in any case, making it arbitrary and insincere to object to their use here. To worry that such invocation is *too late* is to worry that such invocation is already implicit in the skeptical challenge as we have set it up in this paper. Let us conclude by addressing each of these worries in turn.

Concerning the first worry, two points must be made. First of all, to repeat, this paper's argument does *not* stem from any general distrust of modal contrivances or their analytic employment in other areas of philosophy (e.g., analyses of agency, laws, probability or causation). We need not be concerned with arguing that counterfactual explications of philosophical notions are generally suspect. Our concern here, rather, is to argue that such explication is *particularly*

suspect in the context of CCC arguments for reasons that do not apply to such explication across the board.

As a case in point, let us consider just one example of counterfactual explication that is *not* similarly suspect: counterfactual theories of causation. On such accounts, for one event to cause another means that, had the first failed to occur, so would have the second (Lewis 1973c). As we ordinarily talk about knowledge, the invocation of specific modal metrics through the selective privileging of similarity relations plays as great a role here as it does in the case of the CCC strategists' explication of knowledge; for we can hardly maintain that *any* event, but for which an effect would not have occurred, is one of that effect's causes *if* we claim to use "cause" in anything like its ordinary sense. A trauma victim's head injury is caused by said trauma, we want to say, rather than by the victim's possession of a head. Some respects of similarity must be chosen as particularly salient to the exclusion or belittlement of others if causes are to be distinguished from mere background conditions. Examples such as this might suggest that we are wrong to isolate the CCC strategists' use of modal metrics as particularly pernicious; for how can we criticize this employment of such metrics while possibly continuing to allow their use to analyze notions such as causation?

There are, however, crucial differences between the problems of choosing salient similarity criteria as they arise in the above causation case, and as they arise in the case of CCC. For one thing, our concerns about the latter stem from specific worries about the threat of circularity, not from the simple fact that such criteria must be ultimately, when all is said and done, *chosen*. For another thing, it is an option to argue in the causation case that there really is no fundamental difference between "causes" and "background conditions," and that this distinction is purely nominal or instrumental. This is a far from unpalatable position when it is taken to imply, not that causation is fictitious, but that *both* prompters and background conditions are causes if *either* is. The distinction between prompters and background conditions that is present in our ordinary causal talk, on such an account, merely reflects varying explanatory interests without threatening the very existence of the relation we take to be at issue. But a similar move can hardly be made by the CCC

strategists. For them to concede the arbitrariness of the similarity criteria underlying their modal metric is to concede their inability to isolate a knowledge relation to which their investigations are directed at all, even a contextually equivocal knowledge relation on which epistemic demands vary along a single scale of severity.

Such is my response to the reservation that the CCC strategists' invocation of modal machinery might be *too selective*. What about the fear that it is too late? The concern here is that such an invocation may be already implicit in the skeptical challenge of ultimate concern to us; for don't we need modal talk to even set up the skeptical challenge described at the beginning of this paper? Certainly we have articulated skepticism in precisely such terms, with our talk of the inability to rule out alternative possible ways the world might be. Skepticism, by our description, arises for the following reason. We cannot know that some radically extraordinary alternative skeptical scenario doesn't apply because we cannot rule out the possibility that it does. But, how can we now eschew the use of modal discourse by the enemies of skepticism without invoking a double standard which excuses our own use of it?

Once again, the first thing to note is that our objection is not to the use of modal metrics *per se*, but only to their use in programmatic settings that threaten circularity. The second thing to note is that it is far from clear that modal machinery is really even required to articulate the skeptical challenge in the first place. Thus, even if our objection *were* to the use of modal machinery *per se*, we could still maintain that such use is harmless as a useful heuristic device, and is illegitimate only when it cannot be effectively discharged once this heuristic value is exhausted. I suggest that modal language can be discharged from our original skeptical position in a way that it cannot be discharged from CCC arguments. In the case of the former, modal language is defensible as little more than a picturesque way of highlighting the underdetermination of knowledge claims. To say that we cannot rule out skeptical scenarios is to point out that the corpus of available evidence does not conclusively endorse our ordinary worldview. But this is merely to say that deductive links between various sets of sentences are lacking, an idea for which no modal explication is necessary (given a construal of consequence itself in proof-theoretic, rather than

model-theoretic, terms). For DeRose, however, counterfactual invocations are hardly dispensable. A modal metric of objective and ineliminable similarity relations is an essential background feature of his account. Consequently, it is indispensable to his efforts to upholster contextualism in a way that allows it to remain distinct from a simple naturalistic background externalism. Finally, note that this is the case irrespective of any issues regarding the commitment or non-commitment to modal realism. It would indeed be perverse to try to rescue one small class of relations (e.g., knowledge, justification) between people and their environment by positing a countless universe of modal denizens. But our objection is not this. It is broader than this; our concern is with the objective status of similarity standards for comparing possible worlds irrespective of how literally we may take “possible worlds” discourse itself.

Appeals to counterfactual machinery, I have argued, are of little use in grounding efforts to contextualize externalist approaches. Neither, we should add, is such talk needed to ground externalism itself. The suspicion that it might prove necessary is fueled by our schematic characterization of externalism as the position that our knowing *p* requires that we stand in appropriate reliabilist relations to the fact of *p*, relations that make beliefs the results of reliable belief-forming processes. For how, we might wonder, is the probabilistic aspect of reliabilism to be characterized, except in ineliminably counterfactual terms?

Once again, however, there is a ready response to this problem. We can invoke finite frequency or propensity interpretations of probability ascription as alternatives. In finite frequency accounts a belief-forming process is reliable if it offers an adequate likelihood of true beliefs, which likelihood is identified as the number of actual true beliefs formed across a suitable number of employments of the belief-forming process, divided by the total number of employments (Venn 1876). Propensity accounts treat probability as a physical tendency or disposition to produce a specific kind of physical outcome (Popper 1959). Even though neither of these options is without problems (e.g., single-instance probability attribution in the case of the first; an undeniable air of mystery in the case of the second), neither is ultimately more problematic than infinite limiting

relative frequency accounts, which *do* invoke modal machinery in ways that leave empiricism far behind.

## References

- Baumann, Peter (2005). 'Varieties of Contextualism: Standards and Descriptions.' *Grazer Philosophische Studien* 69: 229-245.
- Bennett, Jonathan (1984). 'Review of *Counterfactuals*.' *The Canadian Journal of Philosophy* 4: 381-402.
- Bennett, Jonathan (2003). *A Philosophical Guide to Counterfactuals*. Oxford: Oxford University Press.
- Bird, A. (2005). 'The Dispositionalist Conception of Laws.' *Foundations of Science* 10: 353-370.
- Bowie, G (1979). 'The Similarity Approach to Counterfactuals: Some Problems.' *Noûs* 13: 477-98.
- DeRose, Keith (1995). 'Solving the Skeptical Problem.' *Philosophical Review*, vol. 104, no.1: 1-52.
- DeRose, Keith (1999). 'Contextualism: An Explanation and Defense.' *The Blackwell Guide to Epistemology*: 187-205.
- Fine, Kit (1975). 'Review of *Counterfactuals*,' *Mind* 84: 451-8.
- Goldman, Alvin (1967). 'A Causal Theory of Knowing.' *Journal of Philosophy*, v. 64, 357-372.
- Goldman, Alvin (1976). 'Discrimination and Perceptual Knowledge.' *Journal of Philosophy*, vol. 73: 771-791.
- Gould, Stephen Jay (1981). *The Mismeasure of Man*, New York: W.W. Norton & Co.
- Greco, John and Sosa, Ernest, ed. (1999). *The Blackwell Guide to Epistemology*. Malden, Mass.: Blackwell Publishers.
- Jackson, Frank (1977). 'A Causal Theory of Counterfactuals,' *Australasian Journal of Philosophy* 55: 3-21.
- Krasner, Daniel and Heller, Mark (1994). 'The Miracle of Counterfactuals: Counterexamples to Lewis's World Ordering,' *Philosophical Studies* 76: 27-43.
- Lewis, David (1973a). *Counterfactuals*. Cambridge, MA: Harvard University Press.
- Lewis, David (1973b). 'Counterfactuals and Comparative Possibility.' *Journal of Philosophical Logic* 2: 418-46.
- Lewis, David (1973c). 'Causation.' *Journal of Philosophy* 70, 556-67.
- Lewis, David (1979). 'Counterfactual Dependence and Time's Arrow.' *Noûs*, Vol. 13, No. 4: 455-476.
- Lewis, David (1986). *On the Plurality of Worlds*. Malden, Mass.: Blackwell Publishers.
- Neta, R. (2003). 'Contextualism and the Problem of the External World.' *Philosophy and Phenomenological Research* 66: 1-31.
- Nozick, Robert (1981). *Philosophical Explanations*. Cambridge: Harvard University Press: 167-228.
- Popper, Karl (1959). 'The Propensity Interpretation of Probability.' *British Journal of the Philosophy of Science* 10: 25-42.
- Rescher, N., ed. (1970) *Studies in Logical Theory* (Oxford: Blackwell).
- Richards, T. (1975). 'The Worlds of David Lewis,' *Australasian Journal of Philosophy* 53: 105-118.
- Slote, Michael (1978). 'Time in Counterfactuals,' *Philosophical Review* 87: 3-27.
- Stalnaker, Robert (1970). 'A Theory of Conditionals,' *Studies in Logical Theory*, 98-112.
- Venn, J., (1876). *The Logic of Chance*, 2nd ed., Macmillan and Co; reprinted, New York, 1962.
- Williams, Michael (2001). 'Contextualism, Externalism and Epistemic Standards.' *Philosophical Studies* 103: 1-23.
- Williams, Michael (2004). 'Knowledge, Reflection and Skeptical Hypotheses.' *Erkenntnis* 61: 315-343.

---

<sup>i</sup> DeRose assures me in correspondence, by the way, that this is essentially the view that he continues to hold. Although he does occasionally employ a "double-safety" locution to accommodate the demand that knowledge be subject to upset by nearby worlds in

---

which p is true but the subject disbelieves it (as well as by worlds in which the subject believes p even though p is false), he regards this to be “a new name” for what he takes to be “the same old picture.”

<sup>ii</sup> This question has been posed by others, sometimes in passing (e.g., Neta 2003, 16, fn. 51) and sometimes more systematically (e.g., Baumann 2005), but toward argumentative ends somewhat different from those pursued in the present paper.

<sup>iii</sup> The classic articulation of this objection is offered by Kit Fine (1975). Similar objections can be found in Bennett (1974), Bowie (1979), Jackson (1979), Richards (1979) and Slote (1978).

<sup>iv</sup> See Krasner and Heller (1994) for a number of clever counterexamples to Lewis’ principles.

<sup>v</sup> Bauman (2005, 233) offers a comparative analysis in similar terms, as well as an alternative circularity objection, but by way of a different route and somewhat broader brush strokes than those I aspire to employ here.