

SKEPTICISM, CONTEXTUALISM, EXTERNALISM AND MODALITY

RON WILBURN

University of Nevada, Las Vegas

Abstract

In this paper, I argue for the following claims. Contextualist strategies to tame or localize epistemic skepticism are hopeless if contextualist factors are construed internalistically. However, because efforts to contextualize externalism via subjunctive conditional analysis court circularity, it is only on an internalistic interpretation that contextualist strategies can even be motivated. While these claims do not give us an argument for skepticism, they do give us an argument that contextualism, as such, is not likely to provide us with an argument against skepticism.

1. Introduction

Let's take external world skepticism (henceforth, *skepticism*) to be the position that worldly knowledge is unattainable because various skeptical possibilities cannot be effectively countered, where such skeptical possibilities involve *malign genies*, brains in vats and other sundry mechanisms for sabotaging our ordinary knowledge-acquisition procedures. Epistemic contextualism (henceforth, *contextualism*), by contrast, is the position that standards of knowledge appropriately vary with assertional context. So understood, contextualism is often invoked as a strategy for circumventing skepticism by way of the following reasoning. Because counterpossibility relevance varies with one's assertional situation, no single criterion for reasonable doubt can be reasonably mandated across all epistemic contexts. Thus, even if apodictic requirements are appropriate to the distinctive context of skeptical inquiry, they do not exercise revisionary authority over the many other contexts of epistemic inquiry characteristic of everyday life and science. Skepticism, on this account, is a localized consequence of considerations that apply only within the particular and idiosyncratic context of philosophical inquiry. As such, it is unable to undermine the knowledge claims that properly apply within ordinary claim-making contexts.¹

In what follows, I do three things. In Section 2, I argue that contextualism, to have any hope as an anti-skeptical strategy, must be construed externalistically.

Principia, 10(2) (2006), pp. 171–87. Published by NEL — Epistemology and Logic Research Group, Federal University of Santa Catarina (UFSC), Brazil.

In Section 3, I argue that contextualism, when construed externalistically, is untenable as a position distinct from externalism itself. In Section 4, I address some objections and replies to this second claim, particularly those that invoke subjective conditional analysis in an attempt to contextualize externalism through the use of counterfactual resources. Since internalistic and externalist construals exhaust all possible readings of contextualism, I conclude that contextualism, as such, offers no effective strategy against the skeptic.

2. Why Must Contextualism be Construed Externalistically to Serve Anti-Skeptical Ends?

Let's begin by examining the mechanics through which investigatory context supposedly acts to segregate the negative consequences of skeptical reasoning from the positive claim-making procedures of naturalistic epistemic practice. We would do well here to start with David Lewis, who endeavors to articulate in painstaking detail the contextualist principles responsible for delimiting the scope and range of counterpossibility relevance (Lewis 1996). Even though Lewis' declared starting presumption, in this connection, is the infallibility of knowledge, he denies that context-invariant ignorance follows from this assumption. Why? Because, despite the fact that one only knows that P iff P holds in every relevant possibility left uneliminated by S 's evidence, relevant ranges of counterpossibilities shift with investigative context. What we know thus varies with investigative context because what counts as infallibility varies with such context (Lewis 1996: 222ff).

Consider just three of Lewis' contextualist prescriptions in this regard. The "rule of attention" dictates that counterpossibilities are relevant only as they happen to be those to which we attend (Lewis 1996: 230). The "rule of belief" dictates that no counterpossibility may be ignored if it is one that the cognizer correctly or incorrectly accepts (Lewis 1996: 226). The "rule of conservatism" tells us that we may safely accept the presuppositions of our conversational community (Lewis 1996: 230). Such rules are subject to mutual conflict, of course. But this is a complication that need not concern us here.

The role of such rules is to track conversational occurrences that are jointly sufficient to induce shifts in our epistemic standards. Philosophical reflection, on Lewis' account, creates an investigative context in which, because there are no constraints on which counterpossibilities we may invoke, there is no effective limit

to which our epistemic standards must converge. On this account, the skeptic's reasoning is indeed legitimate, but only within the specialized and segregated context of skeptical inquiry. It cannot be generalized to threaten all of our worldly knowledge claims.

The central question we need to ask about this and similar accounts has been posed before. Why should we take such contextualist principles to enjoy truly normative force? (Williams 2001: 14) To answer this question we need to decide what we take such principles to really describe? Are they merely common constraints on conversational practice? And if so, why should we take them to have ultimate authority over what we do and do not know? Why should we take the mere citation of facts concerning the conditions under which various counterpossibilities are ordinarily entertained as relevant to the task of answering a skeptic whose opening dialectical move is the claim that ordinary practice blithely ignores his own distinctions and insights?

Suppose we take rules of attention, belief, conservatism and the like to function in the above-described manner, as canonical regulations of ordinary conversational practice. What else must we assume if we are to non-arbitrarily endow them with normative force and consequent anti-skeptical import? If truth is regarded as the sole or primary goal of knowledge, then features of context can only bear on the correctness of one's claim to know that p by affecting one's ability to register or recognize the truth of p . Thus, given the starting assumption that the sole or primary goal of knowledge is truth, our task becomes that of linking these rules in the right way to the world they concern, as it is this linkage that stands to render relevant features of this world epistemically accessible.² But, how is this to be done?

The internalist must forge this connection through the intermediary of perceptible evidence. However, the traditional skeptic's central suspicion is that this cannot be accomplished. The reason why is a matter of historical record. To restrict oneself to internalistic evidence is to restrict oneself to one's own phenomenal and intentional states, the contents of one's phenomenal and notional worlds. But, of course, such evidence radically underdetermines the naïve physical theory defining common sense, as well as the more complex physical theory defining natural science. Thus, suppose we that construe the rules of attention, belief and conservatism in an internalistic fashion. The skeptical challenge immediately re-emerges in renewed form. Why should introspectable facts concerning one's attentive focus, prior convictions or presuppositions about community consensus have any bearing at all upon which counterpossibilities deserve addressing before

we proclaim on the attainment or non-attainment of an objective knowledge relation? Barring an extreme externalism concerning reference itself, whereupon contextual factors are themselves internalistically construed, we remain within the very notional world from which the Cartesian skeptic finds no escape. Internalist contextualism will not save us from the skeptic. Rather than segregating the skeptic's negative conclusions about knowledge to a "distinctively philosophical context of inquiry," it sets up the very context from which these conclusions seem so inevitable.

3. Why Contextualism is Significantly Unmotivated When Construed Externalistically

If the preceding is correct, then internalist contextualism falls short as a ground for response to the skeptic because it fails to provide an effective link between contextual factors and the truth conditions of the claims we purport to know, a link which might promise to explain how the accessibility of the latter could vary as a function of the former. What about externalist contextualism?

For the externalist, such a link between contextual factors and the truth conditions of the claims we purport to know is ultimately beside the point. Why? Suppose, for a moment, that one is an externalist about knowledge, where this position is fleshed out in either causal or reliabilist terms. Leaving aside incidental epicycles and qualifications, externalists take one's knowing that *p* to consist in one's standing in appropriate causal or reliabilist relations to the fact that *p*. What could contextualist considerations add to such an account? That is, what form could we allow a distinctively externalist contextualism to take? Certainly such a position couldn't invoke the likes of Lewis' principles of attention and belief, since it is these very principles whose normative status we are trying to validate. As noted above, to so validate these principles' normative status we would need to defend their systematic responsiveness to relevant external states of affairs. But what could insure such responsiveness other than the very causal or reliabilist relations that the externalist, *qua* externalist, is already prone to invoke? Consequently, on the externalist's account, for any claim *p*, one would *already* either know or not know that *p* as a function of those very occurrent causal or reliabilist connections with one's environment. Additional contextualist considerations could appeal to nothing more the very causal or reliabilist relations that characterize the pre-existing stance of externalism. Or, more pre-

cisely, additional contextualist considerations could consist of nothing more than those very causal or reliabilist relations that characterize the pre-existing stance of externalism *if* these contextualist considerations are defensibly construable as reflective of anything deeper than mere facts about linguistic practice. This is an all-important qualification given the aim of defending the systematic responsiveness of contextualist criteria to relevant external states of affairs, since it is precisely this that we need to do if we are to defensibly construe these criteria as reflecting anything deeper than incidental features of how we happen to talk.

The moral of this is immediate: *contextualism has nothing to add to externalism*. It is only as an internalist position that contextualism can even be motivated as a *distinct* position that does not collapse into a background externalism. This is because it is only as an internalist position that contextualism has anything potentially (though not actually, as the skeptic claims) to contribute to the reinstatement of knowledge. Contextualist considerations from the standpoint of externalism are necessarily redundant.

4. Contextualizing Externalism, and the Circularity Objection

The conclusion of the preceding section is worth pausing over. For there have certainly been a number of contextualists who have proffered their positions from staunchly externalist standpoints. Consider DeRose, who aims, at least in his early work, to contextualize Nozick's truth-tracking account, itself a paradigmatically externalist theory. On Nozick's telling, one knows that p only if one is sensitive to p 's truth value, where said sensitivity is analyzed in counterfactual terms. One is sensitive to the truth value of p only if one would not believe p in the nearest possible world in which p is false. On DeRose's account, this sensitivity requirement is then contextualized. In ordinary contexts, as opposed to skeptical ones, we assess the claim to know that p by advancing our perspective, accordion-like, to the nearest possible world in which p is false. 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 radical non-ordinary possibility (involving *malign genies* or the

like). But in skeptical contexts things are very different. In these contexts, (1) the nearest possible world in which Descartes' belief that he is not being deceived by a genie is one in which he is being so deceived, and (2) it is part of the genie's MO to render herself undetectable in principle, jointly imply (3) one 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, 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 one in which the genie is not at his post, as it were, rendering the conditional "If Moore were being deceived by a *malign genie*, then Moore wouldn't believe that he weren't" trivially true by virtue of its false antecedent (DeRose 1995).

This certainly purports to be an account 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 nature, to render them distinct from the more naturalistic and empiricist criteria from which externalist contextualism must otherwise proceed. This is because the counterfactual contextual factors that DeRose invokes here go far beyond the simpler naturalistic and empiricist causal or reliabilist relations described above in Sec. 3. To maintain that contextualism is essentially redundant *vis-à-vis* externalism, therefore, it is necessary to argue against the tenability of this theory and all other theories relevantly similar to it.

Fortunately, at least to my purposes, we can do this by focusing on the role that modal machinery plays in such accounts. The reservation I recommend here does *not* stem from any broad distrust of modal contrivances. Irrespective of whether or not counterfactual explications of philosophical distinctions are generally suspect, they are clearly suspect here. The reason for this concerns the status of the modal metric that must presumably be in place whenever we endeavor, *to anti-skeptical ends*, to zone neighborhoods of possible worlds into those nearer and farther out. Such modal metrics are prone to objective indefensibility, as they must be selectively cherry-picked in advance to support particular anti-skeptical agendas.

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 a handless occupant of some radical skeptical scenario? In the case of many skeptical scenarios (e.g., involving disembodied minds) this might be arguable, but it is not clearly arguable in the case of

all. Imagine, for instance, that our *malign genie* takes the form of the mad scientist of recent philosophical vintage, maliciously poking at one's brain, which we may take to have been envatted mere moments before. The difference between this and the actual world need not involve a radical divergence from our own world's fundamental background laws, but could easily accrue as a function of mere variations in initial conditions. Thus, our question quickly becomes the following. 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 comparison could itself ever settle the matter. To determine our relative proximity to various possible worlds, we would need to consult standards up to the task of objectively dictating which respects of similarity and difference matter more than others, where such *objectivity* would require 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 our ultimate goal in offering DeRose-type accounts is to argue that it is the radical skeptical scenarios that define those worlds most metaphysically distant from our own. But to do this in an informative way, 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 the standard epistemic claim-making procedures of each world fail to be veridical or knowledge-acquiring in the other. Using these terms, we can now state our challenge more succinctly. To cogently employ DeRose's strategy, we must argue that metaphysical distance tracks epistemic distance, but without implicitly presupposing that it does. Again, this is necessary if DeRose is to informatively maintain e.g., that the nearest possible world in which one does not have hands is a world in which one does not believe that one does. Can this be done?

Note here that we cannot simply maintain that any world in which one loses one hands to an industrial accident (henceforth, an *IA world*) is closer than any world in which one's brain is envatted (henceforth, a *BIV world*) to actuality in that it varies less in its basic constitutive atomic attributions of properties to space-time regions. At first blush this might look promising. Certainly when

we consider alternative scenarios *at the moment* one's knowledge of one's own handedness becomes at issue, BIV worlds look to be much more broadly different than do IA worlds from actuality. But this comparison is far from conclusive, since the differences between these worlds and our own cannot date merely from this single selective moment in time. This fact renders the task of comparison much more complicated and problematic. Why?

Suppose that the worlds to which we are comparing our own are all law governed (i.e., "deterministic" in the very general sense that their future states are determined in *either* a unique or probabilistic fashion).³ On such a presupposition, our contemplated "single selective moment in time" must itself have a causal or determinative history, making these worlds' unfortunate divergence from actuality the result of either a difference in background laws or a difference in preceding initial conditions. These are differences that we are obliged to tally and ponder in the course of contrasting worlds. Thus, the observation that BIV brain worlds differ from actuality more than do IA worlds at the moment one's knowledge becomes at issue, *in itself*, tells us nothing. For it to imply our own greater metaphysical distance from BIV worlds (relative to our distance from IA 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 this can be assumed?* If we concede that numerous different networks of causal antecedents can 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.

Alternatively, suppose that these worlds to which we compare our own are *not* all law governed. Such a move might seem to aid the contextualist's choice of modal metric by way of the following strategy. We might endeavor to ensure a minimum of difference between our contemplated IA world and actuality by simply stipulating that the former suffers a fleeting lapse in determinism or change of law only at the moment the accident occurs. This, however, will also not give us what we want, at least if we assume, quite reasonably, that the actual world is, in fact, law governed (i.e., uniquely *or* probabilistically deterministic). Even though such a non-deterministic world might leave pre-accident history identical to that of the actual world, it would still leave us with the task of showing that its variance from actuality, *vis-à-vis* its indeterminism, is of less significance than whatever differences might distinguish various radical skeptical scenario worlds from our own.⁴ In the case of many skeptical scenarios (e.g., involving disembod-

ied minds) this is clearly arguable, but it is not clearly arguable in the case of BIV scenarios. The difference between BIV worlds and actuality need not involve any radical divergence from our own world's fundamental background laws. It could easily accrue as a mere function of different initial conditions.

So far, my claim is a modest one: DeRose gives us little reason to make us think that IA worlds are more metaphysically distant from actuality than are BIV worlds. This claim, then, motivates the driving suspicion of this paper: DeRose's choice of modal metric is not objective. He cherry-picks his modal metric precisely for its compliance with the epistemic metric which he aims to use this modal metric to ground. In itself, however, mere suspicion is of little dialectical consequence. Our real question must be the following: Do we have specific reason to think that convincing arguments for the objectivity of DeRose's modal metric cannot be provided?

I would suggest that we have two reasons to think that no such arguments wait in the wings, the first one theoretical and the second one more concrete. The theoretical reason is the following: *The knowledge relation is only one amongst many that collectively constitute our position within the world we occupy.* Thus, we must ask, why should we think that those worlds in which our belief-forming mechanisms remain reliable are the ones that are metaphysically closest to our own? Wouldn't such an assumption privilege epistemic similarity between possible worlds over all the many other possible respects in which such worlds may be metaphysically similar? In zoning neighborhoods of possible worlds into those nearer and farther out, we must be wary of privileging the similarity salience of precisely those relations that we ordinarily take to be conducive to knowledge itself. We should be wary of criteriologically privileging this salience for the purposes of zoning modal neighborhoods, and then pretending to have grounded our epistemic metric in an independent account of what the modal ordinances dictate.

More concretely and specifically, however, I would suggest that there is *particular* cause to be suspicious of DeRose's counterfactual reliance. For, in the case of DeRose's reliance, we have special reason to worry that the threat of circularity loiters constantly in the background. Consider 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 in which 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 DeRose's modal metric or else epistemic relevance is itself regarded as largely *constitutive* of this metric. Let's focus on the second option, as 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 seems to be a harmless qualification required to address Nozick's "grandmother objection," and its subsequent requirement that beliefs track truth by way of the same method across different 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.) However, on deeper consideration, this proposal should give us pause. The reason why lies in the radical context dependency of epistemic relevance itself, 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 truth-track 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 specifically to 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 visually inspects her grandson. On such a description of belief forming method, BIV scenarios are dismissed from consideration from the start (given the eyeless condition of granny's envatted brain), whereas on more minimalist descriptions of the relevant belief forming method (e.g. consultation with and assessment of relevant experience) they are not. The point is that the very sense of salience that DeRose invokes 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 he must avoid.

5. A Slingshot Circularity Counter-Argument: Too Selective and Too Late?

Here a number of potential objections arise, however. In particular, might not our reservations concerning modal machinery seem both too selective and too late? Don't we need modal discourse 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 our inability to rule out general alternative possible ways the world might be. Skepticism, on this characterization, arises for the following reason. We cannot know that some radically non-ordinary alternative skeptical scenario doesn't obtain because we are unable to 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 with which to excuse our own use of it?

There are two things to note here. First, to repeat, the argument of this paper does *not* turn upon any general distrust of modal contrivances or with their employment in other areas of philosophy (e.g., analyses of agency, laws, probability or causation). Our concern has been to argue that such explication is *particularly* suspect in the context of DeRose's antiskeptical argument because of the influence that his programmatic goals are likely to exercise on the modal similarity criteria he employs to pursue them. It would have to be shown that there is equal reason to suspect that modal discourse in other areas of inquiry court the same danger.⁵

The second thing to note in response to the allegation that our objection is too selective and too late is that it is far from clear that modal machinery is, in fact, ultimately required to articulate the skeptic's challenge in the first place. Thus, even if our objection *were* to the use of modal machinery *per se*, we might still maintain that such use is harmless as a useful heuristic device and 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 articulation of the skeptic's position in a way that it cannot be discharged from DeRose-style arguments. In the case of the former, it 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 the fact that our ordinary world view is not conclusively endorsed by the corpus of available evidence. 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 ineliminatable 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. Note that this is the case irrespective of any issues that might arise concerning one's commitment or non-commitment to modal realism. It would indeed be perverse to try to rescue one small class of relation (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, as our concern is with the objective

status of similarity standards for comparing possible worlds, irrespective of how literally we take “possible worlds” discourse itself.

Appeals to counterfactual machinery, we 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 so necessary is fueled by our schematic characterization of externalism as the position that one’s knowing that p requires that one stand in appropriate reliabilist relations to the fact that 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 in non-counterfactual terms?

Once again, however, there is a ready response to this challenge. We can invoke finite frequency or propensity interpretations of probability ascription as alternatives. On finite frequency accounts, a belief-forming process is reliable if it offers an adequate likelihood of true beliefs, where this likelihood is identified as the number of actual true beliefs formed across a suitable number of employments of said belief-forming process, divided by said total number of employments (Venn 1876). On propensity accounts, probability is treated 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.

6. Conclusions

In this paper, we have argued for the following claims. Contextualist strategies to tame or localize epistemic skepticism are hopeless if contextualist factors are construed internalistically. However, because efforts to contextualize externalism via subjunctive conditional analysis court circularity, it is only on an internalistic interpretation that contextualist strategies can even be motivated. While these claims do not give us an argument for skepticism, they do give us an argument that contextualism, as such, is unlikely to provide us with an argument against skepticism.

References

- Cohen, S. 1988. "How to Be a Fallibilist." *Philosophical Perspectives* 2: 91–123.
- DeRose, K. 1995. "Solving the Skeptical Paradox." *Philosophical Review* 104(1): 1–52.
- . 1999. "Contextualism: An Explanation and Defense." In Greco & Sosa 1999, pp. 187–205.
- Greco, J. & Sosa, E. (eds.) 1999. *The Blackwell Guide to Epistemology*. Malden, Mass.: Blackwell Publishers.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Heller, M. 1999. "The Proper Role for Contextualism in an Anti-Luck Epistemology." *Philosophical Perspectives* 13: 115–29.
- Lewis, D. K. 1973a. *Counterfactuals*. Cambridge, MA: Harvard University Press.
- . 1973b. "Counterfactuals and Comparative Possibility." *Journal of Philosophical Logic* 2: 418–46.
- . 1979. "Counterfactual Dependence and Time's Arrow." *Noûs* 13(4): 455–76.
- . 1986. *On the Plurality of Worlds*. Malden, Mass.: Blackwell Publishers.
- . 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74: 549–67.
- Popper, K. 1959. "The Propensity Interpretation of Probability." *British Journal of the Philosophy of Science* 10: 25–42.
- Sosa, E. 1999. "How to Defeat Opposition to Moore." *Philosophical Perspectives* 13 (Epistemology): 141–53.
- Venn, J. 1876. *The Logic of Chance*. 2nd ed. New York: Macmillan and Co; reprinted 1962.
- Williams, M. 2001. "Contextualism, Externalism and Epistemic Standards." *Philosophical Studies* 103: 1–23.

Keywords

Skepticism, contextualism, DeRose, externalism, knowledge

Ron Wilburn
 Department of Philosophy
 University of Nevada
 Box 455028
 Las Vegas, NV 89121
 USA
 rojobn@unlv.nevada.edu

Resumo

Neste artigo, argumentamos em favor das seguintes afirmações: as estratégias contextualistas para amenizar ou tornar localizado o ceticismo epistêmico são irrealizáveis se os fatores contextualistas forem interpretados internalisticamente. Contudo, uma vez que os esforços para contextualizar o externalismo por meio de uma análise através de condicionais subjuntivos correm o risco de circularidade, é apenas em uma interpretação internalista que as estratégias contextualistas podem sequer ser motivadas. Ainda que essas afirmações não nos dêem um argumento em favor do ceticismo, dão-nos um argumento de que o contextualismo, como tal, provavelmente não será capaz de fornecer um argumento contra o ceticismo.

Palavras-chave

Ceticismo, contextualismo, DeRose, externalismo, conhecimento

Notes

¹ Discussions of contextualism suffer from an obvious problem of branding. That is, the view now comes in so many fragrances and flavors that one must step very lightly in any attempt to characterize a generic contextualist position. One chief branding distinction is that between *attributor* and *subject* contextualism (or “subject sensitive invariantism”), as we decide whose context is to determine the fortunes of knowledge claims (DeRose 1999: 190–1, Cohen 1987; Hawthorne 2004, ch. 4). 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 it were, 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). While 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 recognizable “knowledge relation” is even in play at all). In any case, as will become clear, the “contextualism” of this paper can be narrowly construed as the classic view proffered by Keith DeRose in “Solving the Skeptical Paradox”, as well as any others relevantly similar to it (DeRose 1995).

² Of course, one might *not* take truth to be the sole or primary object of knowledge. This is certainly an option, but not one that I think need give us pause in the context of this *particular* debate. For what the skeptic can say at this point is that the real work of the contextualist's critique is still waiting to be done by an argument showing that the point of knowledge is best construed in other than truth-regarding terms. To show this, one would need to argue that truth is not a legitimate aim of knowledge, so that there is no appropriate sense of loss to be felt at the fact that we have failed to effectively target it. Or, alternatively, and ultimately equivalently, one might continue to treat truth as the sole or primary end of knowledge, but only as "truth" has been reconstructed in a deflated and epistemized form. In either case, the task at hand becomes that of showing that the objects our knowledge claims regard fail to possess completely mind-independent characters. And this is an argument that contextualists, *as such*, do not offer.

³ Probabilistic "determination," so conceived, comes in degrees. For present purposes, however, we need not imagine that the determinative specificity of the effect of the past on the future is any looser than it is in the actual world, where merely statistical determination affords us an immense amount of prediction and control.

⁴ Much ink has been spilled, of course, in an effort to stipulate rules up to the task of ranking criteria of similarity and dissimilarity between possible worlds, most notably by Lewis (1973a, 1973b, 1979, 1986). This is not an issue that we can deal with in the present paper. Suffice it to say that there is good reason to doubt Lewis' ranking criteria. See especially Krasner and Heller (1994) for some clever counter-examples.

⁵ As a case in point, consider counterfactual theories of causation. On such accounts, for one event to cause another is for it to be the case 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 would thus seem to play as great a role here as it does in the case of the DeRose's 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. From examples like this, it might seem that we are wrong to isolate DeRose's use of modal metrics as particularly pernicious. For how can we criticize the employment of such metrics while allowing their use to analyze other notions, particularly notions which, like "causation," we may suspect we require in order to understand the skeptic's challenge. After all, the worry that some radically non-ordinary alternative skeptical scenario might obtain is, in very large part, the fear that the causal order may be not as we suppose.

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 DeRose's strategy. 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. The worry is not that the contextualist's metaphysical similarity metric must be chosen on the basis of some or other criteria, but that this choice is guided by the very same epistemological similarity metric that the contextualist hopes to ground. 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 the distinction is a purely nominal or instrumental one. 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 DeRose. For him to concede the arbitrariness of the similarity criteria underlying his modal metric, he must, in effect, concede his inability to objectively isolate a knowledge relation to which his investigations are directed at all, even a contextually equivocal knowledge relation on which epistemic demands vary along a single scale of severity. This is not an option for DeRose. However, it may be an option for Williams, for whom there is literally no single relation of knowledge with which our concern remains constant across different disciplinary contexts (Williams 2004: 193).