Transglobal Evidentialism-Reliabilism

David Henderson · Terry Horgan · Matjaž Potrč

Abstract We propose an approach to epistemic justification that incorporates elements of both reliabilism and evidentialism, while also transforming these elements in significant ways. After briefly describing and motivating the non-standard version of reliabilism that Henderson and Horgan call "transglobal" reliabilism, we harness some of Henderson and Horgan's conceptual machinery to provide a non-reliabilist account of propositional justification (i.e., evidential support). We then invoke this account, together with the notion of a transglobally reliable belief-forming process, to give an account of doxastic justification.

Keywords Evidentialism · Reliabilism · Justification · Transglobal reliability

1 Introduction

Evidentialism and reliabilism, two prominent philosophical approaches to epistemic justification, typically are taken to be incompatible in a number of important ways. Evidentialism gives pride of place to the idea of contentful evidential-support relations among propositions (relations of "propositional justification"). In order for an agent's belief that p to be justified (to possess "doxastic justification"), according to evidentialism, two constitutive requirements are (i) that the agent possesses evidence whose content provides propositional justification for p, and (ii) that the agent believes that p because of this evidence. Although doxastically justified beliefs typically are produced by reliable belief-forming processes, this etiological fact is not constitutive of doxastic justification. A belief qualifies as justified not because of

D. Henderson
University of Nebraska-Lincoln, Lincoln, NE, USA

T. Horgan (✉)
Department of Philosophy, University of Arizona, Tucson, AZ 85721-0027, USA
e-mail: thorgan@email.arizona.edu

M. Potrč
University of Ljubljana, Ljubljana, Slovenia

the reliability of the process that produced it, but rather because it is based on good evidence. (Reliable belief-forming processes are reliable because they track evidential-support relations.) Evidential support, construed as a matter of contentful relations among propositions, is treated as fundamental and constitutive within the evidentialist construal of justification, whereas reliability of belief-forming processes is derivative and non-constitutive. (One influential articulation of evidentialism is Feldman and Conee 1985).

Reliabilism, on the other hand, gives pride of place to belief-forming processes, and to doxastic justification. A constitutive requirement for an agent's belief that p to be justified, according to reliabilism, is that the belief has been produced by a reliable belief-forming process. Possession of good evidence for a proposition p, according to reliabilism, is to be understood counterfactually in terms of reliable processes: roughly, it is a matter of possessing information I and a reliable belief-forming process Π such that (i) if Π were to be applied to I, then the result would be a belief that p, and (ii) it is not the case that one possesses further information I* and a reliable belief-forming process Ω (perhaps Π itself) such that, if Ω were to be applied to I+I*, then Ω either would produce a belief that not-p or would produce a state of suspended judgment about p. (Alvin Goldman calls this "ex ante" justification, whereas an actually-held belief that is doxastically justified has "ex post" justification.) Reliability of belief-forming processes is treated as fundamental within the reliabilist perspective, whereas evidential support is treated as a derivative attribute that is characterizable counterfactually in terms of process reliability.[1] (The *locus classicus* for reliabilism is Goldman 1976; see also Goldman 1986).

In this paper we will propose an approach to epistemic justification that incorporates elements of both reliabilism and evidentialism, while also transforming these elements in significant ways. The result will be an account of epistemic justification that largely dissolves the tensions between these two positions.

In Section 1 we will briefly describe and motivate the nonstandard version of reliabilism recently advocated by Henderson and Horgan (2006, forthcoming a, b), which they call *transglobal* reliabilism. One striking feature of this position is the extent to which it smoothly accommodates various problem cases that have been put forward as counterexamples to standard versions of reliabilism. In Section 2 we will decouple the notion of reliability from some of the other conceptual machinery deployed by transglobal reliabilism, and we will harness this residual conceptual machinery to provide a non-reliabilist account of propositional justification, i.e., of evidential support. We call this account *transglobal evidentialism*. In Section 3 we will propose a rapprochement between transglobal evidentialism and transglobal reliabilism; we call the resulting position *transglobal evidentialism-reliabilism*. On one hand, certain key claims about epistemic justification that transglobal reliabilism treats as conceptually grounded necessary truths will turn out to have that status after all. On the other hand, the same will be true for transglobal evidentialism. But on a

---

[1] In this paper we concentrate upon the opposition between reliabilism and evidentialism about epistemic justification. We are aware though that there are opposed conceptual pairs regarding justification that follow other parameters. One of these is the distinction between externalism and internalism, which is in the vicinity of our target conceptual dualism but it is not identical with it.

third hand (so to speak), one tenet of standard reliabilism will be eschewed—viz., the idea that process reliability is conceptually and constitutively more fundamental, vis-à-vis epistemic justification, than is evidential support.

2 Transglobal Reliabilism

Transglobal reliabilism is a species of reliabilism. It holds that a necessary condition for being objectively justified in holding a given belief is that this belief be fixed by way of processes that are reliable—where reliability, as always, must be understood as relative to some reference class of environments. Transglobal reliabilism is distinguished by its conception of the relevant reference class of environments. The more familiar versions of reliabilism treat the relevant form of reliability as reliability relative to the agent's own actual global environment. For transglobal reliabilism, on the other hand, the relevant form of reliability is *transglobal* reliability (as Henderson and Horgan call it: 2006; see also Potrč et al., forthcoming) —i.e., reliability relative to the set of *experientially possible global environments*. A possible global environment is *experientially possible* just in case it is compatible with having experiences of roughly the character of those that agents actually have. The actual global environment is but one among a diversity of experientially possible global environments.[2]

Some experientially possible global environments would be extremely epistemically inhospitable—there would be few if any globally reliable processes to be had there. Demon-infested global environments, and those featuring agents as envatted brains, would be cases in point. The actual global environment (which, we take it, is demon free) is a moderately hospitable experientially possible global environment; there would seem to be both more and less hospitable global environments. In yet more hospitable experientially possible global environments, there would be fewer sources of error than there are in the actual global environment. Transglobal reliabilism, in a somewhat simplified nutshell, is the view that it is constitutively required for objective epistemic justification that a belief be the product of processes that are transglobally reliable—reliable relative to the class of experientially possible global environments.

It should be noted that this approach does not require, in order for a process to be transglobally reliable, that the process be globally reliable in all the experientially possible global environments. Rather, and again somewhat crudely, the approach requires that the belief-fixing process be reliable in a *wide range* of such global environments. The operative notion of a "wide range" can be understood thus: the

---

[2] The class of experientially possible global environments is characterized by reference to the actual global environment: each experientially possible global environment is one in which there could be agents whose experiences are roughly like the experiences had by actual agents in the actual global environment. Once the experientially possible global environments are thus specified, we take it, there is no need to invoke the distinction, central in two-dimensional modal semantics (e.g., Chalmers 1996, Jackson 1998), between two ways of modally evaluating a proposition: considering a possible world "as actual" vs. considering it "as counterfactual." In a given experientially possible global environment, a given belief-forming process is either globally reliable (borderline cases aside) or it isn't, *punkt*. So the distinction between global and transglobal reliability seems orthogonal to the two kinds of modality that figure in two-dimensional modal semantics.

exceptional global environments, the ones in a complementary "narrow range" where the process is not globally reliable, all have a distinctive common feature—viz., that *the given belief-forming process is globally unreliable, within the given environment, only because of highly specific, highly coordinated, aspects of that environment that are systematically outside the agent's cognitive ken*. (Demon-infested global environments, and those featuring agents as envatted brains, provide examples.) Vary those aspects a bit in such a way that the resulting global environment is still experientially possible, and you get a global environment outside the narrow range of exceptions and back inside the wide range—i.e., a global environment in which the process is once again reliable.

Consider the parallel matter involving what might be termed *local* versus global reliability. Local reliability is reliability relative to some local environment that an agent encounters within that agent's actual global environment. Within the global environment, there are relatively epistemically hostile and relatively hospitable local environments. Were there a fake-barn region within the agent's global environment, it would constitute a relatively inhospitable local environment vis-à-vis the formation of perceptually based beliefs about "barnhood". Obviously, within our actual global environment, there are local environments that are epistemically relatively inhospitable (thus, the fog of war, the fog of American politics, the fog of fundamentalism, the fog of hypoxia at altitude in a white-out, and so on). There are also relatively hospitable local environments (presumably one's own kitchen is relatively hospitable for perception of everyday objects, with its lack of camouflage, with its paucity of fakes, with the prevalence of familiar objects, and with the presence of oxygen). It is clear that a process can be globally reliable while failing to be locally reliable in some inhospitable local environments. Similarly, a process can be transglobally reliable while failing to be globally reliable in all experientially possible global environments.

Henderson and Horgan (2006, forthcoming a, b) motivate transglobal reliabilism in two stages—a procedure we will also follow in our summary presentation here. The first stage involves motivating a refinement of classical reliabilism, in part by appeal to various thought-experimental scenarios; they call the refined position *neoclassical* reliabilism. The second stage involves yet further refinement in the same general spirit, in part by appeal to yet further thought-experimental scenarios; the upshot is transglobal reliabilism.

2.1 Stage 1: Neoclassical (or Global) Reliabilism

A generic form of reliabilism has become fairly standard. It emphasizes reliability within the global environment that the epistemic agent *actually occupies*—as opposed to various non-actual global environments, such as that of an envatted brain. Adherents of this general position—call it *classical reliabilism*—commonly have not distinguished sharply, or have not made much of, the distinction between local reliability and global reliability. We now consider scenarios suggesting that, were the choice between just these two forms of reliability, then reliabilist accounts of objective justification would best focus on global reliability. Thus suggested the refined reliabilism may be called *neoclassical reliabilism* or *global reliabilism.*

Suppose that Athena and Fortuna are driving from New York City to Memphis. In rural West Virginia, they drive through a county in which there happen to be numerous extremely realistic-looking fake barns within view—although neither of them has any inkling of this fact or any reason to suspect it. As it happens, in this local area all the real barns are yellow, and none of the fake barns or any other buildings are yellow. Again, they have no information to this effect. As they drive past a saliently presented yellow building, Athena, who has had reasonable experience with barns, gets a clear look at it, and on the basis of its barn-like visual appearance, she judges it to be a barn.

Fortuna gets only a very brief glimpse of the building. She saw her first barn just yesterday, elsewhere, and it happened to be yellow. She judges, on the basis of the briefly-glimpsed building's yellow color, that it is a barn—even though she did not get a good look at it, and was, thus, unable to discern any features that are generally distinctive of barns as opposed to other kinds of buildings. It's not that she has a *general* belief that all and only yellow buildings are barns, or that all barns are yellow; she has never formed any such belief. Also, it's not that she has a *general* tendency to inductively extrapolate from old cases to new ones in a hastily-generalizing way, a tendency she is exhibiting here. Rather, it just happens that *in the present circumstances*, a psychological process is present within her that takes as input both the brief glimpse of a yellow building and the yellow-barn memory, and generates as output the barn-belief about the briefly glimpsed object.

So Athena and Fortuna each form the belief that the building is a barn. And indeed it is.

First consider Athena. One's strong inclination is to say that her belief that the building is a barn is extremely well justified. After all, she has excellent perceptual evidence for the belief, and she formed the belief using a process of perceptual barn-categorization that generally works extremely well. On the other hand, as is commonly stressed in discussions of fake-barn scenarios, one is also strongly inclined to say that she does not *know* that the building is a barn, because her belief has been produced by a belief-forming process that happens to be *unreliable in this specific local environment*. Given all those fake barns around, any of which she would have mistakenly taken to be a real barn, the truth of her belief is too much a matter of epistemic luck for the belief to count as knowledge.

Now consider Fortuna. Here one has a very strong inclination to say that she *lacks* objective justification for her belief that the building glimpsed is a barn. This seems so despite the fact that this belief was produced by a belief-forming process that happens to be reliable in the local vicinity she then occupies. The trouble is that *this reliability is itself too fortuitous*, too much a matter of epistemic luck, for the belief to count as justified. And as a consequence, the truth of her belief is also too much a matter of luck for the belief to count as *knowledge*—the reliability of the belief-forming process notwithstanding.

If one takes these intuitive verdicts to be correct, would this mean that we have here a case in which one person's belief (Athena's) is justified even though it is *not* produced by a reliable belief-forming process, and in which another person's belief (Fortuna's) fails to be justified even though it *is* produced by a reliable belief-forming process? That would be too flat-footed a moral to draw; one cannot flatly say that the relevant process is reliable, or that it is not reliable.

Athena employs a barn-categorizing belief-forming process that is indeed highly reliable in one respect: it is globally reliable. Recall that she has had reasonable experience with barns, and it is plausible to suppose that she has thereby come to have a reliable perceptual ability to discriminate barns from non-barns. (In the world at large, there are few fake barns. Humans with reasonable experience typically do become reliable perceptual judges with respect to such generally non-tricky perceptual matters involving common middle-sized physical objects.) But her barn-categorizing process is not locally reliable, because of all those fake barns in the vicinity. Intuitively, the strong objective justification her belief possesses does indeed involve production by a reliable belief-forming process—viz., a globally reliable one. On the other hand, intuitively, the fact that this strongly justified belief nonetheless fails to qualify as knowledge also involves a failure with respect to reliability of the belief-forming process—viz., a failure to be locally reliable. Both dimensions of reliability thus figure importantly: the presence of global reliability seems to figure in her belief being well justified, whereas the absence of local reliability seems to figure in that belief nevertheless failing to qualify as a case of knowledge. (At any rate, this is a prima facie reasonable suggestion for making sense of one's judgments here).

Fortuna is a converse case. Her belief is produced by a barn-categorization process that is locally reliable (for reasons beyond her cognitive ken) but is globally unreliable. One judges that Fortuna's belief fails to be well justified. Intuitively, what appears to figure importantly in its lacking objective justification is the fact that the belief-forming process is not globally reliable. Given this lack of global reliability, the fact that the process happens to be locally reliable strikes one intuitively as a *lucky accident*, epistemically speaking; and for this reason, the belief also does not count as a case of knowledge.

Call a disposition *robust* if it obtains relative to a fairly wide reference class of potential circumstances, situations, or environments. A robust disposition is one whose possession does not depend heavily upon certain unusual or atypical features that are highly specific to the particular circumstance or environment which the possessor of the disposition might happen to occupy; i.e., the disposition does not obtain only relative to a narrow reference class of environments in which those particular features happen to be present.

Now consider reliability, with robustness in mind. The tendency to produce (mostly) true beliefs must be understood as relative to a reference class of actual or possible environments. If a process is globally reliable, it has this tendency with respect to the wide reference class comprising the potential local environments to which an agent might be exposed (or which the agent might inhabit) within that agent's global environment. This is to have reliability in a reasonably robust fashion. For a process to be globally reliable is for its reliability not to depend heavily upon certain unusual or atypical features that are highly specific to the particular circumstance or environment which the possessor of the disposition might happen to occupy; its reliability then does not obtain only relative to a narrow reference class of environments in which those particular features happen to be present. This said, a process can be globally reliable while failing to be locally reliable with respect to some local environment afforded by the global environment (as illustrated by Athena's processes), and a process can be locally reliable without being globally reliable (as illustrated by Fortuna's processes). When a process is locally, but not globally reliable,

that process's reliability does depend heavily upon certain unusual or atypical features that are highly specific to the particular circumstance or environment which the possessor of the disposition occupies. Such merely local reliability is non-robust, because it involves a narrow reference class: local environments with highly specific features exhibited by the particular local environment that an agent happens to be in.

Here is a useful way of thinking about or intuitively gauging the robustness of the reliability of a process or item, a way of doing so that will prove useful as discussion proceeds. Start with a process or item that is reliable in a given environment (for now, think of a particular local environment). Then think of varying the environment in various ways. Robustness of reliability is a matter of not being highly dependent on particular environmental conditions—a matter of accommodating more rather than less environmental variation.

Return now to Athena and Fortuna. From the perspective of classical reliabilism, the apparent lesson of one's intuitive verdicts about them is that in order for one's belief to be doxastically justified, the belief must be the product of a belief-forming process that possesses a *suitably robust* form of reliability. How should one understand this idea of a suitably robustly reliable process? One might think, with an eye on the Athena/Fortuna cases, that global reliability per se is just the ticket. But a moment's reflection reveals that this proposal is too crude as it stands. Consider, for instance, Diana and Delia—another pair of characters in the fake-barn environment. Like Athena, both Diana and Delia have had reasonable experience with barns, and have trained-up, generally serviceable, barn perception processes. Suppose that Diana and Delia both have full knowledge that the particular local region in question is full of very realistic fake barns, and of the location of this region. They both have read all about it in the local morning paper, having spent the night at a hotel along the interstate located just a few miles from the region in question. They now find themselves driving through that peculiar local region (and they know this). It remains true of each of them that her perceptual barn-categorization process is *globally* reliable; after all, the global environment is pretty much as it was before, so those processes remain globally reliable. As they drive along, they see a barn-looking object off in the distance. Suppose now that Diana refrains from forming the belief that it is a barn—because of the very real and salient epistemic possibility that it is a realistic-looking fake barn. And suppose that Delia forms the belief that the object is a barn—on the basis of its barnish appearance, and in spite of what she read this morning in the local paper. (Delia fails to bring to bear her knowledge that there are lots of fake barns in the vicinity.) Now, Delia's belief that the object is barn is indeed produced by a globally reliable belief-forming process, viz., perceptual-appearance-based barn-categorization. (In this respect, Delia resembles Athena.) Nevertheless, Delia most certainly is *not* justified in coming to believe of the barn-looking structure off in the field, on the basis of its visual appearance, that it is a barn. Rather, in these circumstances she should refrain from forming that belief in that way—as does Diana. After all, *they each have excellent reason to believe—indeed, they each know full well—that, in their current local environment, the reliability of this belief-forming process is compromised (it is locally unreliable).* Considered of themselves, those perceptual belief-forming processes remain globally reliable, are here locally unreliable, but in this case would not give rise to objectively justified beliefs (unlike in the case of Athena).

Likewise, consider the case of Elena. Suppose she knows, concerning that particular local area, both (a) that there are lots of barn facsimiles and also (b) that all genuine barns and no other structures of any kind (including both other kinds of buildings and barn facsimiles) are bright yellow. Finding herself in that very location (and knowing it), she catches a fleeting glance of a bright yellow structure, and promptly forms the belief that it's a barn. Of itself, such a process of barn-categorization, on the basis of glimpsed yellow color, would be globally *un*-reliable. But, now, unlike the case of Fortuna, although Elena formed her belief on the basis of a globally unreliable process (viz., classifying briefly glimpsed structures as barns on the basis of their bright yellow color), one judges that this belief of hers *is* justified, even so. Such processes here are locally reliable, globally unreliable, and yet here feature in the production of objectively justified beliefs (unlike in the case of Fortuna). What is going on?

An important general point emerges from the examples of Diana, Delia, and Elena: agents like humans deploy various belief-forming processes in ways that are *holistically integrated* within the agent's overall cognitive architecture. Very frequently, such processes are employed not in isolation, but rather under the modulating influence of various other or wider cognitive processes that are coupled to them and are poised to modulate them if and when certain forms of information become available to those wider processes. The given belief-forming process is under the *modulational control* of these associated processes. Such control can make for a selective application of the process or a selective inhibition, or can otherwise tailor its application to aspects of those local environments about which information is had—and thereby can enhance its reliability as so tailored. In principle, a whole host of different conditioning or modulating relations might be epistemically important. The wider processes might give rise to a narrower process—designing it or otherwise selecting or spawning it. They might trigger the conditioned process in ways that are fitting, or thought to be appropriate. (One's perceptual processes for spotting large, nonhuman, omnivores are primed when walking in known grizzly territory.) They might inhibit it—making for a more selective use of the process. (One's perceptual processes for spotting large, nonhuman, omnivores are given less free reign when walking about the city.) In triggering or inhibiting the process, the wider processes draw upon information that the wider processes generate or possess. All such kinds of modulation can and should be found among normal human cognitive agents.

From a reliabilist point of view, the presence and efficacy of suitable control processes, within the cognitive system viewed as an integrated whole, thus emerges as extremely important in relation to objective justification. What count, though, as *suitable* control processes? What counts as being under suitable modulational control? This looks to be a question that cannot be answered in any great detail from the armchair. What ultimately makes for satisfactory, epistemically competent, belief formation cannot be settled in detail without taking into account which processes can be tractably employed (perhaps innately, perhaps via learning) by cognitive systems of the relevant sort—e.g., adult human cognizers. So, what counts as suitable modulational control depends in part upon the cognitive architecture of the given kind of cognitive agent, and upon the agent's susceptibility (by virtue of its cognitive architecture) for learning or internalizing various kinds of controlling processes that might not be innate to it.

With the notion of modulation by control processes in hand, and also the notion of suitable modulational control, a natural-looking refinement of classical reliabilism now suggests itself, as a way of accommodating cases like those of Athena, Fortuna, Diana, Delia, and Elena. Instead of just saying (as does classical reliabilism) that in order to be doxastically justified a belief must be produced by a reliable belief-forming process, the refined approach affirms the following thesis, which Henderson and Horgan call both "neoclassical reliabilism" and "global reliabilism":

– (GR) In order to be doxastically justified, a belief must be produced by a process that is *globally reliable under suitable modulational control*.

This refinement of classical reliabilism is still very much in the spirit of the classical reliabilist approach to justification.

2.2 Stage 2: Transglobal Reliabilism

Reliability of a process, as noted already, is a feature that obtains relative to a reference class. The wider the reference class, the more robust the associated form of reliability. Global reliability is more robust than mere local reliability, and for this reason it is evidently important for epistemic justification.

Moreover, again as noted already, global reliability of a process is a matter of the process working well within a *wide range* of potential local environments. Thus, a process can be globally reliable while yet failing to work well within certain peculiar local environments—as illustrated by the case of Athena in fake-barn country. The operative notion of working well "within a wide range" of local environments is essentially this: when the process *fails* to work well within some local environment, this failure depends heavily upon certain unusual or atypical features that are highly specific to that particular environment—for instance, the local presence, unbeknownst to the cognitive agent, of numerous highly realistic-looking fake barns.

Even more robust than global reliability is *transglobal* reliability—i.e., reliability relative to the reference class of experientially possible global environments. Thus, in light of the evident importance of robustness, the possibility arises that transglobal reliability is even more important for epistemic justification than global reliability. Moreover, since transglobal reliability is a matter of global reliability in a *wide range* of experientially possible global environments, perhaps there are certain experientially possible global environments in which epistemic agents employ belief-forming processes that are transglobally reliable but yet fail to be globally reliable. (These would be global environments in which the processes fail to work well because of certain unusual or atypical features specific to these environments, as compared to experientially possible global environments in general.) And perhaps the beliefs thus formed would be epistemically justified despite this lack of global reliability—much as Athena's barn-belief is epistemically justified despite a lack of *local* reliability.

Concrete examples are readily given. Just as there are thought-experimental situations suggesting that global reliability of the belief-forming process is more important than local reliability as a constitutive requirement for doxastic justification (and that local reliability is not a constitutive requirement), there are still further scenarios suggesting that *transglobal* reliability is more important yet (and that neither local nor global reliability is a constitutive requirement). Henderson and

Horgan describe a range of cases to make the point, but for present purposes we will focus exclusively on a scenario that is a variant of one that was set out by Lehrer and Cohen (1985) and was labeled the "new evil demon problem" for reliabilism by Sosa (1991). The problem is not about knowledge (as is the original evil demon problem), but is about justification. It involves scenarios in which an experiencer has beliefs that intuitively count as well justified, but that reliabilism counts as unjustified.

In keeping with philosophical mythology, suppose that there is an evil demon—malicious and very powerful—out to deceive the agent. Seeking to epistemically defeat the agent at every turn, the demon provides the agent with appearances or experiences that are compellingly consistent in the way they indicate an apparent world in which the experiencer is an apparently embodied, apparently perceiving, apparently acting, cognitive agent; these experiences are just as compellingly consistent as are the embodied, en-worlded, experiences of ordinary humans. As the agent undertakes to do things in that world, the demon responds by giving the agent the fitting appearances. But, the world that the agent and the demon inhabit is radically other than the world that the agent is led to believe in and theorize about—thus the deception. Of course, if the demon is really good at deception, then the agent is presented with sufficient epistemic problems to "solve"; these keep the agent engaged in making sense of the seeming-world the agent experiences.

Objectively, there is no hope for an agent in a demon infested environment. There are no globally reliable processes to be had in such an environment. Whatever cognitive processes the agent employs, the demon will counter with input that is fittingly deceptive (that leads to false beliefs) given those processes. (At least this holds for empirical beliefs that the agent might generate. The case of a priori beliefs is not addressed here. Also, we are here assuming that empirical beliefs pervasively of the form "The Deceiver is providing me excellent evidence for the false statement that …" are non-starters as candidates for being objectively justified—even though they happen to be true—since agents in the envisioned scenario have not a shred of evidence for such radically paranoid empirical beliefs).

These days, computers are commonly cited in place of evil demons when constructing skeptical scenarios. Typically, one supposes that the supercomputer has charge of a brain in a vat. Again, the envatted brain has no real chance of arriving at systematic true belief.[3]

---

[3] Cartesian Evil Demons and envatted brains remain as popular as ever as thought-experimental tropes in epistemology—despite the fact that in philosophy of mind, various strongly externalist views about mental intentionality entail either that an envatted brain's beliefs would not be systematically false, or that an envatted brain would not have beliefs (or other intentional mental states) at all. Here we follow the common practice in epistemology of (1) ignoring strongly externalist views about mental content, (2) assuming that the envatted brain has beliefs and other intentional mental states that in a strong way *exactly match* those of an unenvatted human whose brain undergoes exactly similar neurobiological processes, and (3) assuming that the envatted brain's beliefs are systematically false. For one approach to mental intentionality that defends assumptions (2) and (3) and argues that strongly externalist positions in philosophy of mind are seriously mistaken, see Horgan and Tienson (2002) and Horgan et al. (2004). (These papers advocate an exact match in phenomenally constituted *intentionality* between unenvatted and envatted agents with duplicate neural histories; on the other hand, mental *reference* for certain thought-constituents—e.g., constituents that purport to refer to individuals or to natural kinds—constitutively requires the joint interaction of phenomenal intentionality and certain externalistic linkages between agent and environment. On this treatment, many of the envatted brain's beliefs not only are false, but also suffer reference failure because of the lack of suitable internal-external linkages).

Whether plagued by evil demons or envatted, an agent in such an exceedingly inhospitable global environment is doomed to form non-veridical beliefs. There are no reliable processes to be had in such an environment. Whatever cognitive processes the agent employs, the demon or computer will counter with input that is fittingly deceptive (that leads to false beliefs) given those processes. The agent will end with systematically false beliefs.

We turn now to an example of the new evil demon problem, formulated in terms of envatment. Consider envatted counterparts of our five thought-experimental heroines from above; call them Envatted Athena, Envatted Fortuna, and so on. Each of them has lifelong experiences exactly matching those of her non-envatted counterpart. Consider the pertinent barn-beliefs of these folks, formed when they have their pertinent experiences as-of driving from a place called New York City to a place called Memphis. Intuitively, the pertinent barn-beliefs of each envatted experiencer have *exactly the same justification-status* as do the corresponding barn-beliefs of her unenvatted counterpart. Envatted Athena, for example, is very well justified in her perceptual-experience-based belief that there's a barn over yonder—just as much so as is (Unenvatted) Athena. Envatted Fortuna, on the other hand, is not justified in believing (on the basis of a very brief experience as of a yellow structure of some sort or other) that there's a barn over yonder—her belief too has the same epistemic status as does the corresponding belief of (Unenvatted) Fortuna. And so forth.

An adequate account of doxastic justification, we take it, should accommodate these intuitive verdicts about Envatted Athena and Envatted Fortuna. But neoclassical reliabilism evidently cannot do so; this is the new evil demon problem in action. The trouble is that under neoclassical reliabilism, all the pertinent beliefs of each of these experiencers—and virtually all of their other external-world beliefs too—count as entirely unjustified. Envatted Athena's barn-belief, for instance, counts as unjustified because it is produced by a process that is *not globally reliable*—not globally reliable in itself, and not globally reliable under the modulational-control structures to which it is coupled.

Two important morals emerge from such scenarios. First, what is constitutively required for doxastic justification is that the belief be produced by a process with a form of reliability even more robust than global reliability—viz., *transglobal* reliability. Second, a belief can meet this requirement, and hence can be well justified, even if (as in our envatment scenario) the belief-forming process happens to lack global reliability. Both morals are structurally analogous to the morals that emerged earlier, concerning global reliability as more important than local reliability, and concerning the possibility of a belief's being justified despite local unreliability of the process that produced it. (Notice, too, that the neoclassical reliabilist treatment of cases of Athena and Fortuna now gets displaced too: in those cases as well, the constitutively required, highly robust, form of reliability that is present for Athena and absent for Fortuna is not really global reliability, but rather is transglobal reliability.)

The earlier remarks about modulational control still remain applicable, mutatis mutandis. Transglobal reliabilism, then, claims:

- (TGR$_d$) In order to be doxastically justified, a belief must be produced by a belief-forming process that is *transglobally reliable under suitable modulational control*.

(The subscripted 'd' is for 'doxastic', and the feature here invoked is what Henderson and Horgan call *postclassical* reliability.) The cases of Envatted Diana, Envatted Delia,

and Envatted Elena will illustrate the need for the modulational-control clause, in a manner parallel to earlier discussion of their unenvatted counterparts. (Here too the neoclassical treatment of the original Diana, Delia, and Elena gets displaced: what is constitutively required for justification is not global reliability under suitable modulational control, but rather transglobal reliability under suitable modulational control.)

What about propositional justification? Transglobal reliabilism can give an account of that too, by adapting the "counterfactual strategy" described in the second paragraph of this paper. Roughly, at least, the account is this:

- (TGR_p) Proposition B is propositionally justified for an agent A (at a time t) just in case A possesses information I and a postclassically reliable belief-forming process Π such that (i) if Π were to be applied to I, then the result would be a belief that B, and (ii) it is not the case that A possesses further information I* and a postclassically reliable belief-forming process Ω (perhaps Π itself) such that, if Ω were to be applied to I+I*, then Ω either would produce a belief that not-B or would produce a state of suspended judgment about B.

3 Transglobal Evidentialism

Transglobal reliability is reliability in a wide range of experientially possible global environments. Two conceptual components can be distinguished here: on one hand, being reliable (relative to some reference class), and on the other hand, obtaining in a wide range of experientially possible global environments. We will now suggest a way of harnessing the latter component, disconnected from the former, to articulate a non-reliabilist account of propositional justification. The proposed articulation will then be available as the basis for an account of doxastic justification too. Thus, our proposal will lead to a version of evidentialism; we will call it *transglobal* evidentialism.

It is natural to think of evidential support as a relation that is somewhat analogous to entailment. A set of propositions $A_1, ..., A_n$ entails a proposition B just in case B is true in every possible world in which $A_1, ..., A_n$ are all true. By analogy, evidential support can be construed as a suitable weakening of the entailment relation. Roughly and generically, the idea is that B is true in a *wide range* of *relevant* possible worlds in which $A_1, ..., A_n$ are all true.

A natural thought, next, is to construe the pertinent idea of a relevant possible world in terms of the notion of an "experientially possible global environment," as employed by transglobal reliabilism. But this will not do as it stands, because one must distinguish between a possible *global environment* and a possible *world*. Within philosophy and formal semantics, a possible world normally is construed as a totally specific possible *world-history*, a totally specific "way a world might be." (Or else it is some kind of ersatz stand-in for such a thing—e.g., a totally specific *maximal proposition* describing a possible world-history, or a maximal complete *property* instantiable by the cosmos.) What Henderson and Horgan call a "possible global environment," on the other hand, is a world in the more familiar pre-theoretic sense of 'world'—i.e., a possible overall environment capable of being the locus for any of numerous different potential total world-histories. (Or else it is some kind of ersatz stand-in for such a thing).

When discussing a notion like reliability, the notion of a global environment is the more pertinent one. Reliability of a process in a global environment is a matter of that process *generally* working well within that environment—which means that the single global environment is one in which numerous different total world-histories might have occurred. On the other hand, when one sets reliability aside and seeks out an analogue of entailment, the more pertinent notion is that of a possible world. So we will make that adjustment here.

As mentioned earlier, an experientially possible global environment may be characterized as a global environment that is compatible with one's there having experiences of roughly the character of those that agents actually have. In some such environments—e.g., ones infested by Cartesian evil demons or by envatted brains— quotidian belief-forming processes (e.g., processes relying on perceptual experience, or on memory experience, or on inference to the best explanation) are not reliable. (These are global environments in which there is highly specific, very elaborate, and enormously well coordinated behind-the-scenes "jerry rigging" that persistently provides an epistemic agent with ongoing experiences that are highly cohesive and yet utterly nonveridical.) Nonetheless, in a *wide range* of such environments, these processes are globally reliable; i.e., they are transglobally reliable. The operative notion of a "wide range" is the following, with respect to a feature F of an item i (in this case, the feature *global reliability*, instantiable by a given belief-forming process): in any experientially relevant global environment E in which F is not instantiated by i, this is so because of highly specific, highly coordinated, highly interdependent, aspects of E. Vary those aspects a bit in such a way that the resulting global environment is still experientially possible, and you get a global environment outside the narrow range of exceptions and back inside the wide range—i.e., a global environment in which item i once again instantiates feature F.

So the following two ideas now suggest themselves, as steps toward fashioning an account of evidential support. First, retain the idea of a *wide range* of experientially possible global environments, while disconnecting it from reliability of belief-forming processes. (We have just now formulated the idea more abstractly, in terms of an item i instantiating a feature F in a wide range of such global environments.) Second, replace the notion of an experientially possible *global environment* by that of an experientially possible *world*, while retaining the operative notion of experiential possibility. With these two ideas in place, now define the notion of *strong evidential support* this way:

– (E) Propositions A_1, ..., A_n collectively provide strong evidential support for proposition B just in case B is true in a wide range of experientially relevant possible worlds in which A_1, ..., A_n are all true.[4]

---

[4] This definition perhaps needs refining, in order to accommodate so-called "centered propositions." Roughly, a centered proposition is one involving first-person indexical propositional constituents of the kind expressible linguistically by words like 'I', 'now', 'here', and the like. The truth conditions for such a proposition can be modeled as a set of "centered possible worlds," where each member of the set is a possible world with a designated "center" to which the relevant indexical constituents are "anchored." We will not pursue here the question of whether—and if so, how—the proposed definition of strong evidential support should be refined in order to suitably accommodate centered propositions.

In the general case of *non-demonstrative* evidential support, of course, a "total evidence" principle is applicable: even if propositions $A_1, ..., A_n$ collectively provide strong evidential support for proposition B, it may be the case that there is some more inclusive collection of propositions $<A_1, ..., A_n, C_1, ..., C_m>$ that does *not* provide strong evidential support for B. (The propositions $C_1,...C_m$ might include "defeaters" for the evidential support provided by the propositions $A_1, ..., A_n$).

Principle (E) provides a way of thinking about evidential support that makes no appeal to reliability of belief-forming processes. It leads naturally, in turn, to a characterization of *propositional justification* that also eschews appeal to process reliability. Let a *psychological profile* P[A,t], for an agent A at a time t, be a full body of psychological facts about A up through time t—including facts about A's psychological history prior to t, and facts about A's psychological dispositions both past and present (e.g., dispositions to form certain occurrent beliefs under conditions of querying). Let a *psychological duplicate* of an agent A (at t) be an actual or possible agent with a psychological profile that at some moment exactly matches P[A,t]. Exact match can fall short of literal *identity*, because of potential differences such as different referents of indexical terms and concepts, e.g., the essential first-person indexical 'I myself'. But two agents who are psychological duplicates would have psychological profiles that are experientially exactly similar. With these notions in play, we can define propositional justification this way (with 'TGE' being short for 'transglobal evidentialism'):

– (TGE_p) A proposition B is justified, for an agent A at a time t, just in case B is true in a wide range of experientially relevant possible worlds in which there is an agent who is a psychological duplicate of A (at t).[5]

Having thus characterized propositional justification in a way that makes no reference to belief-forming processes or their reliability, we can now build upon definition (TGE_p) to give a transglobal evidentialist characterization of *doxastic* justification:

– (TGE_d) Agent A's belief that B is justified (at t) just in case (i) the proposition B is justified for A at t, and (ii) A believes B (at t) because B is justified for A at t.

There is more to be said, no doubt, by way of further cashing out clause (ii) of (TGE_d); and more *will* be said in the next section. But this clause already makes tolerable sense by itself, intuitively and pre-theoretically—even though the idea it expresses surely deserves further philosophical investigation and elucidation.

Surely there is also more to say, and more that needs saying, both about the idea of an experientially relevant possible world, and about the idea of a "wide range" of such worlds. Each of these notions has received some further elucidation from Henderson and Horgan (see especially Chap. 3 and 5 of Henderson and Horgan forthcoming a), although the ideas deserve yet further exploration. Here too, however, both notions already seem fairly clear pre-theoretically. And we have been

---

[5] The qualification mentioned in note 4 applies again here.

suggesting that they can do some real work: they can underwrite a plausible-looking version of evidentialism—viz., transglobal evidentialism, comprising theses (TGE_p) and (TGE_d).[6]

4 Transglobal Evidentialism-Reliabilism

Return now to the key theses of transglobal reliabilism: (TGR_d), which specifies a key necessary condition for doxastic justification, and (TGR_p), which deploys the "counterfactual strategy" to characterize propositional justification in terms of doxastic justification. What is to be said about these two theses, given transglobal evidentialism?

To begin with, notice that these two theses are both consistent with transglobal evidentialism; they are not in conflict with it. Moreover, the two theses both might have the status of being conceptually grounded necessary truths; that too would not be in conflict with transglobal reliabilism.

Given these two facts, one can envision a potential version of transglobal evidentialism that makes the following claims. First, two theses (TGR_d) and (TGR_p) are indeed true. Second, they are indeed conceptually grounded necessary truths. However, third, the concept of strong evidential support is conceptually more fundamental than is the concept of a postclassically reliable belief-forming process. Likewise, fourth, the concept of propositional justification is conceptually more fundamental than is the concept of doxastic justification. Thus, fifth, the two theses (TGR_d) and (TGR_p) both are *derivative* conceptually grounded necessary truths— i.e., they are logical consequences either of transglobal evidentialism itself, or of transglobal evidentialism together with plausible supplementary assumptions. We will call this position *compatibilist* transglobal evidentialism—the idea being that it is a form of evidentialism that is compatible with theses (TGR_d) and (TGR_p), while also treating process reliability as a decidedly derivative notion insofar as both propositional justification and doxastic justification are concerned.

Compatibilist transglobal evidentialism obviously effects a significant rapprochement between evidentialist and reliabilist ideas. For, it takes on board the central theses of one form of reliabilism (viz., transglobal reliabilism), despite demoting the notion of process reliability to a derivative status conceptually. It is a position worth taking seriously, we suggest.

---

[6] It might be objected that thesis (TGE_p) cannot provide an adequate conceptual handle on the idea of evidential support unless and until one articulates some substantive general principles that govern the notion of evidential support. I.e., it might be held that evidential-support relations must be systematizable by *exceptionless general principles*—perhaps principles that could be expressed as a computer program for manipulating representations of all the intentional content of a given unified body of experience—and that an adequate explication of evidential support must *articulate* these principles. We have two responses. First, the claim that one needs to specify such principles in order to explicate evidential support gets things backwards. Any proposed substantive general principles would need to get assessed for adequacy insofar as they conform to one's pre-theoretic understanding of evidential support as a "makes likely true" relation (even though *some* pre-theoretic intuitions about evidential support might end up being rejected as mistaken). Second, there is good reason both (i) to deny that that there are exceptionless general principles that govern the relation of epistemic support, and (ii) to deny that there need to be such principles in order for human cognition to do reasonably well at forming beliefs in ways that track evidential-support relations. (Compare Potrč 2000, Horgan and Potrč 2006, 2008 and for a wider philosophical perspective Horgan and Potrè forthcoming).

It is not, however, the position we wish to advocate here. Instead, we want to recommend a view that gives postclassical reliability a more central role in the characterization of *doxastic* justification, alongside the transglobal evidentialist treatment of *propositional* justification.

Return again to thesis (TGE$_d$) of transglobal evidentialism, and consider clause (ii) in particular: "[Agent] A believes [proposition] B (at t) because B is [propositionally] justified for A at t." How might this clause be further elucidated? Intuitively, the idea is that A's cognitive system is somehow *using* information I available to A that confers strong propositional justification upon B, and furthermore is using that information in a way that is *sensitive* to the fact that I confers strong justification upon B. One would like a way of plausibly capturing these intuitive ideas. Moreover, certain ways of trying to do so that might initially come to mind are really not very plausible at all, because of the unrealistic and intractable demands they place on human cognition. Specifically, the following picture of justified-belief formation looks utterly unrealistic psychologically: agent A *explicitly represents* all pertinent available information that collectively confers strong evidential support upon proposition B, and *explicitly represents* the way(s) that all this information conspires to render B likely to be true, and then forms a belief that B via *explicit inference* from all this explicitly represented information.

A more plausible approach is to invoke process reliability. The kind of belief-forming process that would fill the bill, with respect to explicating the idea that A's cognitive system is using A's available evidence, qua evidence, in generating the belief B, should be a process that *tracks relations of strong evidential support*. And what kind of process is that? Well, we suggest, a process that is postclassically reliable—i.e., is transglobally reliable under suitable modulational control. Although a process that is merely *globally* reliable might happen to work in a way that does not suitably depend on evidential support—say, because of fortuitous features of the actual global environment that are outside the agent's cognitive ken—transglobal reliability is another matter. Such fortuitous features will not obtain across a *wide range* of experientially possible global environments; and for this reason, postclassical reliability is very tightly connected, conceptually, to evidential support (as explicated by thesis (E) above)—so tightly connected that postclassical reliability constitutes the tracking of evidential support. (For example, imagine a global environment that is very much like our own earthly environment, save for the fact that it was created 10,000 years ago, ex nihilo, by a god who instilled in humans a strong disposition to believe that the global environment was thus created. Exercising that disposition would be a globally reliable process within the given global environment, despite the fact that this process would not be evidence-tracking. Its failure to be evidence-tracking, we are proposing, is constituted by the fact that the process would not work well across a *wide range* of experientially possible global environments).

What we wish to recommend, then, is a position we will call *transglobal evidentialism-reliabilism*. Concerning propositional justification, it embraces thesis TGE$_p$ above, which can also be relabeled:

$$(TGER_p) = (TGE_p)$$

Concerning doxastic justification, however, it replaces (without repudiating) thesis (TGE$_d$) with the following, whose second clause is intended to further elucidate clause (ii) of (TGE$_d$):

– (TGER$_d$) Agent A's belief that B is justified (at t) just in case (i) the proposition B is justified for A at t, and (ii) A's belief that B (at t) is produced by a postclassically reliable process that has operated on information in virtue of which proposition B is justified for A.

The idea of a postclassically reliable process thereby gets an "upgrade to first class," in terms of role. Transglobal evidentialism gives this idea only "tourist class" status, since (TGE$_d$) defines doxastic justification without even mentioning reliable processes—so that postclassical reliability figures only in a thesis that transglobal evidentialism treats as a *derivative byproduct* of (TGE$_d$), viz., thesis (TGR$_d$). But transglobal evidentialism-reliabilism upgrades the conceptual role of transglobal reliability by invoking this notion directly within its more detailed explication of doxastic justification, (TGER$_d$).

It is a question for cognitive science what is the nature of postclassically reliable processes of the kind that figure in clause (ii) of (TGER$_d$). Also a question for cognitive science is just how such processes operate on the evidentially pertinent information available to the agent—although the lately mentioned tractability issues strongly suggest that much of the information needs to be accommodated automatically by the cognitive system, in ways that rely on the morphological structure of the system's cognitive architecture rather than on explicit representations of all this information.[7]

---

[7] It is worth noting that at least two points advanced in this paper reflect an expansive understanding of the evidence possessed by a human epistemic agent, one that diverges from what is commonly suggested by evidentialists, such as Feldman (2004). First, in our formulation (TGE$_p$), we say that evidential justification can be understood in terms of how the proposition in question turns out in those experientially relevant possible worlds where there is a psychological duplicate of the agent whose justification is gauged. Here, being such a duplicate is understood in terms of matching psychological profiles at the time in question, which involve "the full body of psychological facts about [the individual agents] up to the time in question". This strongly suggests that evidential support turns on a wider set of information or contentful states than the sort of access internalist states that are commonly the focus of evidentialists. The suggestion is then made more concrete in the discussion of (TGER$_d$), where the information or content of concern includes that which is accommodated automatically by the agent as a cognitive system, in ways that rely on the morphological structure of the system's cognitive architecture. All this contrasts with the range of alternatives Feldman considers for understanding the evidence possessed by an agent. Feldman (2004: 219) ultimately defends "a restrictive account that limits the evidence a person has at a time to the things the person is thinking of or aware of at that time." He does toy with allowing evidence to include a "feeling of certainty" or confidence—perhaps resulting from a kind of proceduralization—and this does help make some room for background information. But ultimately, it seems, the background information never really counts as evidence for him unless it is both accessible and accessed at the time in question (Feldman: 238–240). From our point of view, this has several problems. First, on the face of it, it does not seem to allow for a needed distinction between cases where the background information makes epistemically appropriate the "feeling of certainty" and those in which it does not. Second, Feldman might be understood as saying that this distinction should be drawn in those cases in which the background is rendered almost foreground—taking the form of "beliefs" that are "unconsciously thought" and that thus are "available" to the agent. For reasons articulated in Henderson and Horgan (2000, forthcoming c), we believe that this seriously underplays the importance of background information.

Such automatically accommodated implicit information is called *morphological content* (Horgan and Tienson 1996, Potrč 1999, 2000, Henderson and Horgan 2000, forthcoming b). Transglobal evidentialism-reliabilism gives pride of place to the notion of strong evidential support, as explicated non-reliabilistically in terms of truth in a wide range of experientially relevant possible worlds (thesis (E) above). That notion is the basis for our recommended account of *propositional* justification, and also figures importantly in the account of *doxastic* justification. Nonetheless, transglobal evidentialism-reliabilism also gives pride of place to postclassical reliability, rather than relegating it to merely derivative status (in the manner of compatibilist transglobal evidentialism). For, postclassical reliability figures centrally in the proposed explication, in clause (ii) of thesis (TGER$_d$), of the idea that agent A believes proposition B *because* B is propositionally justified for A. The upshot, we suggest, is that transglobal evidentialism-reliabilism effects a more thorough, and more satisfying, rapprochement between evidentialist and reliabilist ideas than does compatibilist transglobal evidentialism.

5 Conclusion

The notion of propositional justification is a matter of evidence. Intuitively and pre-theoretically, propositions A_1, ..., A_n provide evidence for proposition B just in case B is *likely true* given A_1, ..., A_n. The reliabilist approach to propositional justification does not directly capture this idea of a "making-likely-true" relation. Although we have suggested here that Goldman's "counterfactual strategy" for characterizing propositional justification in reliabilist terms can be successfully implemented—in the form of the conceptually grounded necessary truth (TGR$_p$), which employs the notion of postclassical reliability—nonetheless we also maintain that (TGR$_p$) is too derivative a truth to really capture the pre-theoretic notion of propositional justification. That notion is better captured via thesis (TGE$_p$), which explicates the idea that proposition B is likely true given propositions A_1, ..., A_n this way: B is true in a wide range of experientially relevant possible worlds in which A_1, ..., A_n are all true. So, as regards propositional justification, transglobal evidentialism looks preferable to transglobal reliabilism; thesis (TGE$_p$) is conceptually more fundamental than is thesis (TGR$_p$), even though the latter is still true according to transglobal evidentialism (and indeed is still a conceptually grounded necessary truth). Transglobal evidentialism-reliabilism, in turn, treats propositional justification the same way as does transglobal evidentialism.

The notion of doxastic justification, pre-theoretically and intuitively, is in part a matter of propositional justification (as pre-theoretically understood); i.e., *part* of what is constitutively required for doxastic justification is that the agent has propositional justification for the proposition believed. Transglobal evidentialism captures this conceptual dimension of doxastic justification, via thesis (TGE$_p$). But doxastic justification is also a matter of believing a proposition *because of* the propositional justification one has for it. Transglobal evidentialism leaves this second element at that, in clause (ii) of thesis (TGE$_d$). But a suitable, psychologically plausible, further explication of this idea is wanted. Transglobal evidentialism-reliabilism provides one: viz., that believing a proposition "because of" one's

propositional justification is a matter of one's belief being caused by a postclassically reliable belief-forming process that employs one's propositionally justificatory information. Transglobal evidentialism-reliabilism thus provides a more thorough, more satisfying, treatment of doxastic justification than does unvarnished transglobal evidentialism. And in so doing, it provides a fundamental conceptual role for the notion of postclassical process-reliability vis-à-vis doxastic justification, even though propositional justification gets treated in a non-reliabilist, transglobal evidentialist, way.

Transglobal evidentialism-reliabilism, then, effects a fairly thorough-going rapprochement between evidentialist ideas (non-reliabilistically construed) and reliabilist ideas. There are important elements of truth in both approaches, but these elements fit together in one cogent package much more thoroughly than has been previously appreciated in epistemology.

References

Chalmers, D. (1996). *The conscious mind*. Oxford: Oxford University Press.

Feldman, R. (2004). Having evidence. In E. Connee & R. Feldman (Eds.), *Evidentialism* (pp. 199–241). Oxford: Oxford University Press.

Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies, 48*, 15–34.

Goldman, A. (1976). What is justified belief. In G. S. Pappas (Ed.), *Justification and knowledge*. Dordrecht: D. Reidel.

Goldman, A. (1986). *Epistemology and cognition*. Cambridge: Harvard University Press.

Henderson, D., & Horgan, T. (2000). Iceberg epistemology. *Philosophy and Phenomenological Research, 61*, 497–535.

Henderson, D., & Horgan, T. (2006). Transglobal reliabilism. *Croatian Journal of Philosophy, 6*, 171–195.

Henderson, D., & Horgan, T. (Forthcoming a). Some ins and outs of transglobal reliabilism. In S. Goldberg (Ed.), *Internalism and externalism in semantics and epistemology*.

Henderson, D., & Horgan, T. (Forthcoming b). *A priori naturalized epistemology: At the interface of cognitive science and conceptual analysis*.

Henderson, D., & Horgan, T. (Forthcoming c). Epistemic virtues and cognitive dispositions. In K. Steuber, G. Damschen, & R. Schnepf (Eds.), *Debating dispositions: Issues in metaphysics, epistemology and philosophy of mind*. Berlin: DeGruyter.

Horgan, T., & Potrč, M. (2006). Particularist semantic normativity. *Acta Analytica, 21*, 45–61.

Horgan, T., & Potrč, M. (2008). Contextual semantics and particularist normativity. In M. Lance, M. Potrč, & V. Strahovnik (Eds.), *Challenging moral particularism* (pp. 123–139). New York: Routledge.

Horgan, T., & Potrč, M. (Forthcoming). *Austere realism: Contextual semantics meets minimal ontology*. Cambridge: MIT Press.

Horgan, T., & Tienson, J. (1996) *Connectionism and the philosophy of psychology*. Cambridge: MIT Press.

Horgan, T., & Tienson, J. (2002). The intentionality of phenomenology and the phenomenology of intentionality. In D. Chalmers (Ed.), *Philosophy of mind: Classical and contemporary readings* (pp. 520–533). Oxford: Oxford University Press.

Horgan, T., Tienson, J., & Graham, G. (2004). Phenomenal intentionality and the brain in a vat. In R. Schantz (Ed.), *The externalist challenge* (pp. 297–317). Berlin: de Gruyter.

Jackson, F. (1998). *From metaphysics to ethics: A defense of conceptual analysis*. Oxford: Oxford University Press.

Lehrer, K., & Cohen, S. (1985). Reliability and justification. *The Monist, 68*, 159–174.

Potrč, M. (1999). Morphological content. *Acta Analytica, 22*, 133–149.

Potrč, M. (2000). Justification having and morphological content. *Acta Analytica, 25*, 151–173.

Potrč, M., Sajama, S., & Strahovnik, V. (Forthcoming). Skeptical justification.

Sosa, E. (1991). Reliabilism and intellectual virtue. In E. Sosa (Ed.), *Knowledge in context* (pp. 131–145). Cambridge: Cambridge University Press.