

Would You Really Rather be Lucky than Good?

*David Henderson, University of Nebraska, and
Terence Horgan, University of Arizona*

1. Introduction

Quine famously clarified his call for a naturalized epistemology, saying:

Naturalization in epistemology does not jettison the normative and settle for the indiscriminate description of ongoing processes. For me normative epistemology is a branch of engineering. It is the technology of truth-seeking... (Quine 1998, 664-5).

In the present quick paper, we discuss the normative force of naturalized epistemology in terms of (a) several forms of reliability (local, global, and what we have termed transglobal reliability), (b) the capacity of humans to modulate and even generate cognitive processes drawing on information provided by wider cognitive processes (what we term modulatory control), and (c) the goal of believing truths or having systems of true beliefs—a goal that believers have by virtue of being believers. We point out that modulatory control, drawing on relevantly reliable wider processes can enhance the reliability of the processes under such control—and that the results are valuable for anyone engaged in the pursuit of true beliefs. Naturalized epistemology is concerned with providing a normative account of belief fixation, oriented by the central epistemic concern of producing systems of true belief, and drawing freely and significantly on empirical results.¹ Naturalized epistemology thus amounts to a paradigmatic form of modulatory control fitting to humans.

Let's begin assembling the leading ideas.

2. Some Thoughts on Reliability.

Reliability in various forms is an epistemic good. Indeed, we think there are valuable forms of reliability that have not been adequately distinguished and appreciated in epistemology.

Reliability, of course, is the tendency to produce (mostly) true beliefs. As all dispositions, it must be understood as relative to a reference class of actual or possible environments. We will call a disposition *robust* if it obtains relative to a fairly wide reference class of potential circumstances, situations, or environments. The idea is that 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.

If a process is *globally* reliable, it has this tendency (to yield true belief) with respect to the wide reference class comprising the potential local environments to which an agent might be exposed 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 to not depend heavily upon certain unusual or atypical features that are highly specific to the particular circumstance or local 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 local environments. This said, for a globally reliable process, there may yet be certain local environments—involving unusual or atypical features—in which the globally reliable process would prove unreliable. A process can be globally reliable without being locally reliable with respect to some local environment afforded by the global environment, and a process can be locally reliable without being globally reliable. When a process is merely *locally* 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: a specific local environment that an agent happens to be in. Local reliability is yet an epistemic good thing, along with global reliability. So also is the relatively greater robustness of reliability.

We have argued elsewhere that there is an epistemically important, highly robust, form of reliability that has been neglected in much contemporary epistemology—what we term *transglobal reliability*. Transglobal reliability is reliability relative to the set of *experientially possible global environments*. A possible global environment is *experientially possible* just in case it is compatible with one's there having experiences of roughly the character of those that agents actually

have.² The actual global environment is but one among a diversity of experientially possible global environments. Some would be extremely epistemically inhospitable. In demon-infested global environments, for example, there would be few if any globally reliable processes to be had. The actual global environment is a moderately hospitable experientially possible global environment. In yet more hospitable experientially possible global environments, there would be fewer sources of error.

While we have argued (Henderson and Horgan 2006, Forthcoming b) that transglobal reliability is pivotal in common judgments employing the concept of *being objectively epistemically justified in a belief*, what we have to say here does not focus specifically on objective epistemic justification, or on transglobal reliability. What we have to say turn more on the normative epistemic concern for various forms of reliability and for relatively robust reliability.

Consider two scenarios in your possible biography:

Early on you are a new and underprepared member of a bomb squad. You are confronted with a mass of wires, and the sense that which one you snip matters very much. You make a guess, and you snip. The timer stops and nothing else happens. “Better lucky than good,” you mumble. But, do you really believe it? Probably not, and you shouldn’t. In forming beliefs about what wires are safe to cut, you would rather be good, systematically good, than merely lucky. Of course the scenario has both an epistemic and a more broadly practical element—but in either case, you would rather be good than lucky. The point can be understood in terms of reliability, of course. Your belief-forming process here (merely guessing) is not a reliable way of forming beliefs about how to safely defuse a bomb—it is not locally, globally, or transglobally reliable—and it should not be employed.³

Suppose that you later take a job at an agency for mental health. On the basis of MMPI scores, you are told to sort patients into those who are psychotic and those who merely neurotic. Your friend the shaman offers to help craft a method. First, he arranges names for the component scales in a circle. Then, twirling an antelope head with grouse feathers in the eye sockets, he assembles a formula this way: if the snout stops when pointing to a given clinical scale, you add that score, if the base stops when pointing to a scale, you subtract that component. Since your friend is 45, you take that as the dividing line, and (unbeknownst to you) you end up with Goldberg’s Rule. It works pretty well; you get promoted. Better lucky than good? Of course not. Most find the whole episode epistemically distasteful, bad form, irresponsible, and unjustifiable—unnecessarily risky. It is not enough to merely deploy a rule or method that happens to yield reliable verdicts. The rule or its use should itself be supported, drawing on the results of well designed research.

The problem here is not simply a failure of reliability. Goldberg’s Rule is as reliable, or more reliable (presumably both globally and locally) than, any

tractable alternative cognitive process of initial diagnostic classification (or so we are told by [Goldberg 1965](#), Dawes 1994, Bishop and Trout 2005a, 2005b). This much—the reliability of the processes you deploy—is an epistemically good thing.

Still, the reliability of a process considered of itself, even its relatively robust reliability, is not the whole story. Somehow, the whole episode strikes one as unacceptably risky. That there might be epistemic misgivings along these lines is not news to reliabilists. Witness Goldman’s (1992) insistence that (acquired) reliable processes (“methods”) should themselves be the result of metareliable processes—processes that reliably generate reliable processes. Sosa (2000) has related concerns. We want to suggest a way of developing these concerns.

3. Package-deal Reliability and Suitable Modulatory Control

The epistemically optimal and epistemically satisfactory levels of reliability for human cognitive agents commonly require a kind of informed modulation of our cognitive processes by information gotten from wider cognitive processes.⁴ Consider perceptual processes. One should confidently render perceptual verdicts only regarding those domains where one’s processes have been trained up. In training one acquires significant ranges of information—much of which then may be employed “automatically” and subdoxastically within the perceptual processes (Henderson and Horgan 2000, Forthcoming a). Training may draw on further perception (on closer inspection, for example), but various inferences and testimony or instruction can play an important role. General knowledge should modulate one’s perceptual processes. For example, information about base rates in one’s region should condition one’s perceptions. The ursine form glimpsed at a distance should be perceived differently when in Glacier National Park on a summer evening and in New York City on Halloween evening.

Consider how experimental design and associated inference are informed by the results of wider processes (including earlier experimental processes). This is reflected in the contemporary demand that claims regarding the effectiveness of a drug be supported by double-blinded placebo-controlled studies.⁵ This demand reflects an important refinement in earlier experimental and inferential practice. The refinement was prompted by studies that had made salient the threats to reliability posed by placebo effects and experimenter bias.

The epistemological concern for the package-deal reliability of yoked processes can be understood in terms of a notion of *suitable modulatory control*. Human cognitive agents deploy various belief-forming processes in

ways that are *holistically integrated* within the agent's overall cognitive architecture. Such processes are employed not in isolation, but rather under the control of various other or wider cognitive processes that condition or modulate them. When these modulating or controlling cognitive processes provide veridical information about the agent's environments, such control or modulation makes for a selective application of the modulated process, and thereby enhance its reliability as so conditioned.⁶

What kinds of conditioning or modulating mechanisms there are in humans, and what kinds humans are *capable* of developing and deploying, are empirical questions, principally within the purview of cognitive science.⁷

Of course, conditioning processes themselves can be reliable or unreliable, and can exhibit various kinds reliability (local, global, and transglobal). A fundamentally important feature that conditioning processes presumably ought to possess is that they be relevantly reliable *themselves*. The following tendencies are clear and significant. Consider a process that is, standing alone, merely *locally* reliable. Its reliability is dependent on certain conditions obtaining in the local environment that an agent happens to occupy for the time.⁸ Suppose that it is under the modulational control of various *globally* reliable processes. They may selectively trigger it on the basis of information about those features of local environments on which its reliability depends, or they may inhibit it in complementary ways based on such information.⁹ They may even have designed the process—again on the basis of information about the global and local environments. In such cases, as the control processes afford the modulation, the controlled or then modulated process tends to become more *globally* reliable *as so controlled*. A process that is under suitable modulational control of a globally reliable process tends to be *globally reliable under suitable modulational control*.

Many aspects of scientific reasoning provide safeguards against various possible pitfalls—think, for example, of the safeguards characteristic of good statistical reasoning. Here, historical mathematical investigations of probability and statistics—work that is itself transglobally reliable, condition one's reasoning about one's data and hypotheses. The reasoning processes that have resulted are *transglobally* reliable. Often, results of past scientific experiment, study, and reasoning are used to inform further investigation—as when transglobally reliably produced information about physical or chemical processes in one's global environment are used to design a device to reliably detect the value of some parameter in some system—and to do so in a way that is not highly dependent on just where that system is located in one's environment. (Think of shielding against spurious events.) This characterizes a common kind of case: one in which a process is designed to be *globally* reliable on the basis of information that is itself the product of processes that are *transglobally* reliable. Were those wider transglobally reliable control processes to have been

at work in a somewhat different global environment, they would likely have yielded different results, which would have informed a process more suited to (more globally reliable in) that different global environment. Again, the controlled process may be only globally reliable, considered of itself, while also being, *transglobally reliable under such suitable modulational control* and ongoing modulation.

In general then, with suitable modulational control in play, with ongoing modulation, the controlled process, considered as under that control, comes to have the reliability status of the controlling process.

4. Reliability, Epistemic Risk and Epistemic Safety

Epistemology revolves around a characteristic end: that of producing systems of true belief. Epistemological evaluations of various stripes are oriented by this central constitutive epistemic end. We want to use such familiar observations in some straightforward ways.

When considering various scenarios and processes, one finds oneself and others objecting or approving—digging in one’s heels or sensing epistemic safe passage. One often finds that it is natural to express one’s approval or misgivings in terms of safety and unacceptable risk. Were one to explore the matter more thoroughly than is possible here (for reasons of time), one would find that judgments about unacceptable risk are keyed to various form of reliability of processes considered both narrowly and as under modulational control.

Most central human endeavors are multifaceted, and subject to various evaluative perspectives. For an analogy, think of the various ways of evaluating the play or players of a game such as baseball. It would be surprising were there some two of three evaluative concepts that lend themselves to all the evaluations one is inclined to make in a domain. It would be surprising were all epistemic evaluations of significance to boil down to whether some agent satisfies the concept of *knowledge*, or *objective epistemic justification*, or *subjective epistemic justification*. This would be stranger than thinking that all evaluations of baseball batters boil down to the application of the concepts of *batting average* and *on-base percentage*. One’s various epistemic evaluations reflect a simple point about this diversity: as an epistemically alive and engaged agent, one cares about various forms of reliability, degrees of robustness of reliability, kinds of modulational control, and resulting forms of epistemic safety. One also cares about the tractability of one’s epistemic projects. Indeed, it seems fair to say that suitable modulational control commonly involves a kind of concern for, or care for, or sensitivity to, the reliability of one’s controlled processes—their transglobal, global, or local reliability, as the case might be.

Now, the (perhaps historical) concern for some of these forms of robust reliability and safety may have resulted in the development of a salient evaluative epistemic concept—one that may have fixated epistemologists. If we (Henderson and Horgan) are right, the concern for transglobal reliability under suitable modulational control is central to the concept of *being objectively justified in holding a given belief*. Such concepts play an important role in one’s self-regulation and an important role in our social regulation of our epistemic community. But to think that epistemic engagement is largely a matter of just applying one or two evaluative concepts (such as objective and subjective justification) is a flatfooted narrowing of the richness of epistemic concern.

5. Finally! The Normative Character of Naturalized Epistemology

The reflections above set the stage for a series of points regarding the epistemic engagement that is commonly termed “naturalized epistemology.”¹⁰ We want to develop these points regarding naturalized epistemology:

1. Naturalized epistemology is a wonderful example of the sort of modulational control that we above identified as a central epistemological concern. It contributes significantly to various forms of reliability, robust reliability, and safety. So understood, naturalized epistemology always represents a desirable and necessary aspect of human epistemic engagement.
2. So engaged, the naturalized epistemologist can for a period be focused on any of various forms of reliability and safety.
3. Sometimes the engaged naturalized epistemologist will clarify and apply one of the familiar concepts that have so exercised philosophers—such as that of *objectively epistemically justified belief*. Such “conceptual analysis” is not foreign to naturalized epistemology. At other times that naturalized epistemologist will be focused on cobbling together processes that serve some situated epistemic purpose. The purpose may be one marked by an entrenched evaluative epistemic concept, or one with less entrenched conceptual honor in the tradition. For example, a naturalized epistemologist might for a period be concerned with what would make for a locally reliable process for a situated set of agents. Here Quine’s idea of an engineering of truth seeking is helpful. Engineers are commonly employed on projects ranging from the highly local (a durable insulation to shield a particular vehicle during reentry into the earth’s atmosphere) and highly general (materials that

might serve some function such as thermal insulation in a wide diversity of environments).

4. To an engaged agent, questioning the normative force of an engineering result, or of an application of a familiar evaluative concept, is likely to seem a nonpressing “philosopher’s question.” It would be rather like asking the engineer designing a reentry vehicle what is the normative force of *being a good insulator* or of the finding that such-and-such a material is a good insulator. It is like asking a baseball manager what is the normative force of *being a good batter*, or the finding that so-and-so is hard to strike out. At one level, the normative force is obvious—conversational rules suggest that the question of the normative force of naturalized epistemology must have a less obvious answer. The normative practice can go forward without an articulate answer to the more abstract question.
5. The philosopher’s questions here are not silly and can be given answers.
6. Contemporary philosophers are correct to be engaged with this full range of cognitive engineering questions. In the past, philosophers have “taken positions” on some of these engineering questions without knowing it—and theirs have commonly been mistaken or flawed positions as a result (see Henderson and Horgan 2000).

Naturalized epistemology is a fine example of the sort of modulatory control that is a central epistemological concern.

For a given process to be under the modulatory control of wider processes is for the agent’s wider cognitive system to be disposed to variously modulate the process as the wider processes provide information. Now, the wider processes can either turn on empirical information in a big way, or turn largely on reflection.¹¹ Clear cases of the latter would be provided by the history of mathematical reflection on probability and statistics. Mathematics has repeatedly yielded results informing experimental and other practice and inference. The ensuing modulation is a form of epistemology that all have welcomed. Modulatory control attentive to empirical information is also thick on the ground—and it is worth considering some diverse cases. As we have noted, perceptual processes within a domain are permeable to much information. For example, such processes are sensitive to information having to do with base rates (or sampling frequencies)—information which results from wider ranges of empirical processes. And again, various forms of investigation—experimental and other—are informed by other empirical investigations. We have already mentioned how the study of prospective drug

regimens has been modulated by information about placebo effects and experimenter bias. Similarly, one who seeks to reconstruct an episode on the basis of witness accounts would do well to study the psychological literature on these accounts and the kinds of suggestion and reconstruction to which they are subject. Various studies of the foibles of human cognitive practice may lead one to seek to structure one's reasoning processes in way that militates against empirically indicated pitfalls—anchoring effects, for example, or overreliance on certain heuristics where these are readily distorting, or the under use of information about base rates (or sampling frequencies).

Naturalized epistemology provides a normative account of belief fixation, drawing freely and significantly (but not exclusively) on empirical results. The just mentioned cases of modulational control drawing on empirical information would seem clear cases. Thus, naturalized epistemology is a kind of modulational control. It is not the only kind, but it is an important kind—it is the systematic kind that draws on empirical results.

Yet, this way of putting the matter may be understood as suggesting that naturalized epistemology is distinct from epistemology in a way that is unhelpful. It should be said that what is commonly understood as naturalized epistemology—normative work drawing on empirical results—and mainline epistemology—drawing significantly on reflection—are, or should be, so entwined as to constitute one endeavor. Thus, naturalized epistemology can, does, and should, draw on results of reflection that are not themselves obviously empirical. For example, studies dealing with how one might improve certain classes of reasoning regarding probabilities (or their frequency parallels) themselves commonly relate the recommended processes to mathematical results. Thus, when [Gigerenzer and Hoffrage \(1995\)](#) discuss “how to improve Bayesian reasoning without instruction” (apparently by couching problems in terms of frequency formats rather than in probability or relative frequency formats), the “improvement” envisioned is itself understood as a matter of sensitivity of information that makes for results conforming to Bayes' theorem. On the other hand, it has come to be increasingly appreciated that fitting normative accounts of how one ought to fix belief must not neglect the real potentialities (capacities and capacities for capacities) of those cognitive systems (typically humans) for which the accounts are intended.

One might quibble. While acknowledging that all naturalized epistemology is a form of modulational control drawing on empirical information, one might insist that not all empirically informed modulational control is naturalized epistemology. If one thinks of philosophy, and of epistemology, as a form of inquiry with a substantial level of generality, then one might insist that some empirically informed modulational control is not sufficiently articulate and general to count. We are not much interested in drawing lines here. We think it best to recall Quine's dicta to the effect that, “Scientific neologism is

itself just linguistic evolution gone self-conscious, as science is self-conscious common sense” (1960, 3). A reasonable suggestion is that naturalized epistemology is particularly general, articulate and self-conscious methodology or methodological practice and engagement. Ongoing empirically-informed modulational control then exhibits a range of cases. There is the inarticulate practice (including empirically informed refinements in perceptual processes). There are more articulate refinements in relatively narrow cognitive practices (such as those dealing specifically with drug testing). There are articulate refinements of general processes (such as abandoning strivings for internalist accessibility, a refinement informed by recent cognitive science). All are very much in the same business. Naturalized epistemology is, it seems, this form of modulational control “gone self-conscious”—and there is not much sense in thinking that there is some sharp or bright line to be drawn here, given the continuities of “the technology of truth seeking.”

Engineering generalities and engineering particulars can be mutually informing and constraining. This provides further reason for resisting the drawing of a bright line between epistemology (as just generality) and methodology or even articulate practice (as just particularity). One can recognize a difference in tendency without supposing a difference in kind.

We have already discussed reasons for thinking that such modulational control is epistemically desirable and to be demanded: to the extent that the empirical processes on which the naturalized epistemologists draw are locally, globally, or transglobally reliable, and correspondingly safe, the resulting modulational control tends to enhance the relevant reliability and safety of the processes under such control. When one thinks that modulational control is humanly feasible and could improve an agent’s epistemic situation—making for greater safety—one is strongly inclined to demand that the suitable processes for modulational control be in place and in play. One finds belief formation epistemically problematic, inappropriate, or seriously flawed when the suitable processes of modulational control are missing or not deployed, and do so even when the process that is deployed without the benefit of such modulation yet happens to be reliable. This is starkly evinced in the case of the shamanic variant on the use of Goldberg’s Rule. Extant psychological evidence indicates that Goldberg’s Rule provides a process for categorizing cases that is reasonably reliable—more reliable than the parallel early-stage categorizations of clinical experts using the same or even richer information. So, one might adopt such a method on the basis of those empirical psychological results. If one were to do so, then not only would one’s categorization process (applying Goldberg’s Rule) considered on its own be a reliable (arguably globally reliable) process,¹² but, additionally, the process would thus be under the modulational control of globally, and even transglobally, reliable psychological methods and results. Nevertheless, applying that same narrow process—

applying Goldberg's Rule—while it happens to be globally reliable—is yet judged seriously problematic and flawed when it results from our agent's reliance on the shamanic coaching (rather than relying on careful psychological investigations and statistical comparisons). Applying a process without suitable, feasible, modulational control is *an epistemic house of cards*—it is unacceptable, unsafe, even when one “lucks out” and the narrow process happens to be relevantly reliable. Modulational control by highly unreliable processes such as shamanic divination would be a case in point.

Still, a serious epistemic defect can arise either by the use of inappropriate modulational control processes—as illustrated by the shaman-based adoption of Goldberg's rule—or by the absence of processes of modulational control where those could readily be had. An interesting case of the latter is diagnosed by Dawes (1994, 38–47). A paradigmatic case of sophisticated and articulate modulational control is provided by the contemporary insistence on randomized experiments with control groups. This demand and practice is a fitting response to compelling experiments that made salient the threat to reliability posed by various biases or foibles in human judgment. But, certain claims for diagnostic capacities and effective treatments are apparently yet arrived at in a fashion that still resists the indicated safeguards: the claims commonly made and accepted to the effect that psychiatrists, clinical psychologists, and clinical social workers have developed capacities for reliable categorization, prediction, and effective intervention. Dawes argues that these claims typically have not themselves resulted from studies with the needed safeguards. Yet, obviously, these are matters regarding which the threat of compromising bias resulting from the ready overuse of various cognitive heuristics should be particularly salient. (Prominent here is the overreliance on the representativeness and availability heuristics—of course there are the additional challenges posed by placebo effects and experimenter bias.) Even so, practitioners and educators have proven highly resistant to the modulational control clearly called for. In fact, when limited studies with the indicated safeguards have been done, these have failed to turn up any evidence for the enhanced reliability of judgment often claimed to result from extensive clinical experience (Dawes 1994, chapters 3 and 4). Yet, the belief-forming processes of practitioners (both regarding individual cases and regarding the reliability of judgments regarding such cases) seem resistant to such modulational control.

There are two faces of a coin here. First, apparently, beliefs about the reliability of certain belief-forming processes have been formed by processes that are themselves lacking in suitable modulational control. The common belief in the reliability of clinical judgment is thus improperly generated—and suitable modulational control would have mandated something on the order of randomized controls and comparisons in the generation of beliefs about the reliability of the judgments of experienced clinicians. Second, given the lim-

ited extant evidence, it seems that suitable modulational control would have called into question the relative reliability of such expert processes, and this should induce modulational control of the epistemic processes for initial diagnosis—yielding a greater use of actuarial processes such as Goldberg’s Rule, and the diminished faith in, and reliance on, clinical judgment.

Some might worry that the demand for suitable modulational control, the demand for epistemically engineering ourselves as truth seekers, is apt to spin out of control—demanding all sorts of cognitively costly control processes, or possible control processes that have not yet been dreamt of. In response, we should emphasize that the notion of *suitable* modulational control must be tied to what is feasible or *tractable for the relevant class of cognitive systems*. (Typically normal adult humans are the concern.)

The naturalized epistemologist can for a period be focused on any of various forms of reliability and safety.

Quine’s idea of an engineering of truth seeking is helpful. Engineers, as engineers, are concerned with the resources for, and design of, safe and reliable (and fast and efficient) systems. First, consider familiar kinds of engineers. Vehicle engineers, are concerned with safe and reliable vehicles. Traffic engineers, with the safe and reliable road systems in which the relevant class of vehicles can be employed at expected levels. Aeronautical engineers, with safe and reliable flight vehicles. Computer engineers, with reliable systems on which one can safely and efficiently undertake certain computational chores. With regard to both the materials for systems, and the design of the systems themselves, engineers can face highly local or highly global problems. Thus, an engineer at work on a vehicle for Antarctic ground transportation is concerned with local reliability and safety. Materials and design are attuned to the demands of an environment featuring extreme cold, high winds, sweeping plains covered by ice and snow, and extended darkness. Some of the materials that the engineer then chooses may have been developed for just such extremes. Others will have been developed their global reliability for certain purposes. A material might be developed to serve as a light yet excellent thermal insulator across temperatures commonly found at the surface of the earth. It might then be selected for use in the vehicle because its global reliability as an insulator finds no exception within the local Antarctic environment. An engineer developing a commercial airliner is concerned with a globally reliable system capable of delivering people and cargo to Juneau or to Jakarta. Of course, some subsystems will be designed for local reliability—for example, the entertainment systems may be designed to work in the climate controlled cabin. Onboard computers also are designed for rather restricted temperature ranges—and the onboard climate control system will need to be globally reli-

able as a consequence. Engineers designing space vehicles may need to design for even greater “global” reliability. Exterior materials must be selected to deal with temperatures encountered in space and upon reentry into planetary atmosphere. That would call for a highly globally reliable thermal insulator, for example. On the other hand, when it comes to getting about on the ground, the vehicle may be tailored to work at a few select landing strips, for example. That is pretty local—as vehicles go. So, when working on materials to be applied across a range of systems and when working on a single system, commonly engineers will have situated concerns for both local and global reliability.

Not surprisingly a similar situated diversity marks one’s epistemic concerns and engagement. One is learning all the time—learning about one’s world and one’s species, about one’s communities, and oneself—with the consequence that one is learning about how one might learn. As one learns, one acquires or refines one’s dispositions—either automatically or with some considered design and effort. Such is the modulational control lauded in this paper. One seeks to fashion oneself and one’s community into an increasingly relevantly reliable, and thus relevantly safe, increasingly efficient, and sometimes increasingly rapid, generator of systematic true beliefs. In effect, then, one is reflectively or unreflectively, engineering oneself (and one’s community) as an epistemic engine.

Let us consider this engaged engineering at its most general. Quine remarks that “science is self-conscious common sense” (1960, 3)—and the remark applies to both scientific understanding and scientific practice. One aspect of scientific practice that self-consciously parallels everyday learning how to better learn is how it draws on information about the global and local environment to tailor investigation and inference. As we have noted, some of this information is about math, statistics, and probabilities. This leads us to incorporate into our epistemic practice safeguards that would add to global reliability in a wide diversity of epistemically possible worlds. This is to engineer for transglobal reliability, and it is a continuing dimension to our epistemic engagement. The scientific practice that benefits from this engagement is relatively transglobally reliable. It produces results having to do with ourselves and our global environment that themselves enjoy a measure of transglobal reliability. These empirical results themselves further inform our ongoing practice as we seek to avoid pitfalls that we come to understand are characteristic of ourselves and our global environment. We may come to appreciate that the subjects of our study are given to placebo effects or effects merely as the result of being studied. We then structure our study and inference to control for these effects. We may find experimenter bias is particularly pronounced in certain contexts, and add corresponding safeguards—such as “blinding” certain investigators regarding who received what treatment. Insofar as these practices or processes

are under the modulational control of relatively transglobally reliable processes, they are themselves package-deal transglobally reliable. Of course, as modulation occurs and one's processes are attuned to one's global environment, their global reliability is augmented. This concern for attuning one's processes, making for increasing global reliability, is a pervasive element of one's epistemic engagement—a pervasive aspect of one's cognitive engineering of oneself and one's dispositions. (Of course, the engineering here can often be automatic—a matter of recognition and reaction that involves little conscious thought.)

One way in which global reliability can be enhanced is by taking advantage of regularities that hold globally. Thus, a detector for certain chemical substances might be globally reliable by virtue of making use of understood chemical regularities that hold universally or very widely within the physical world. Another way in which global reliability is enhanced is by guarding against various local environments where the process would produce false positives or negatives. One then does not use this process or device in the compromising temperature ranges or near to sources of known confounding materials. Here the concern is for both local and global reliability—and the refined processes or dispositions themselves are informed by information that one wants to be globally and transglobally reliable. The device here “embodies” a proceduralized inference, drawing on background information about physical regularities and about the materials there employed. One's restricted use of the device is a refinement of such inferential practices informed by further information.

But, of course, one does not need to wait for a light to come on in some device in order to detect many things in one's environment. Humans are perceptual detectors. One's application of oneself as a perceptual detector is modulated in parallel ways, evincing parallel concerns for local, global, and transglobal reliability.

In some cases, one consciously designs one's investigation and inferences—as in the story of double-blinded, placebo controlled, medical experimentation. In other cases, the process takes shape less self-consciously—as when one simply subjects oneself to courses of experiential training. Here much information is accommodated more or less automatically in shaping one's process. In each case, one finds either a concern for, or inarticulate sensitivity to, various forms of reliability (local, global, and transglobal) and the corresponding epistemic safety provided.

So, epistemic engineering, modulational control of one's cognitive processes in the pursuit of systematic true belief, exhibits concerns for local, global, and transglobal reliability.¹³ This is not quite to say that naturalized epistemologists, *as epistemologists*, are concerned or engaged at each of these levels. Here we encounter the line-drawing suggestion resisted already above. Well,

one might insist on drawing bright lines here, but the cost is high and the payoff low. One thereby makes epistemology a less interesting discipline—as one can then say remarkably little about just how humans should form beliefs. One could say that they should employ processes that are variously reliable under suitable modulational control—but what processes are these? What are they like? These are questions of undeniable epistemological interest. What kinds of modulational control are tractable? On what ranges of information do they draw, and how? Such questions cannot be sorted out without much concrete engagement with cognitive science and psychology—engagement in which one gets down to the grubby matter of just what processes seem to have the various abstract features.

Sometimes the engaged naturalized epistemologist will clarify and apply one of the familiar epistemic concepts; other times the epistemologist will be focused on cobbling together processes that serve some situated epistemic purpose with less entrenched conceptual honor.

Some partisans of naturalized epistemology write as though such inquiry will have no use for concepts such as *objective epistemic justification*, or *knowledge*. They think that an analysis or understanding of such concepts probably provides at most an anthropological or historical insight. Thus, Kitcher (1992) draws on a line of thought from Skyrms (1986) questioning why one should care about satisfying a concept such as *rationality*, or *epistemic justification*. Bishop and Trout evince a similar attitude. But, we think that such misgivings regarding the continuing significance these evaluative epistemological concepts can be answered.

As a sociolinguistic matter, such evaluative concepts are likely to have arisen in earlier engaged epistemic practice—that is, in an ongoing engineering of truth seeking. As a result, they are likely to reflect and to have coalesced around the sorts of engineering concerns that are characteristic of modulational control processes. Thus, these concepts are likely to reflect and revolve around the sorts of engineering concerns that one has today. After all, it should not be difficult to understand why one should continue to care about the local, global, and transglobal reliability of one's belief fixing processes—or about their power and tractability. So, there is reason to expect that contemporaries continue to care, at least about the values motivating or informing these evaluative concepts. They care because they are engaged in the project of finding systematic true beliefs.

It seems that the real misgiving regarding the continuing epistemic significance of concepts such as *objective epistemic justification* or *knowledge* turns on an understanding of these concepts—and one that we think is mistaken. The idea

seems to be that these concepts represent *fossilized understandings* of what processes answer to the epistemic concerns—that they would demand of epistemic agents certain processes that were *hitherto thought to be epistemically reliable and safe*, but which might turn out not to be. Obviously contemporary epistemologists, qua epistemologists, need to be concerned with what *is* epistemically reliable and safe, not with what was once was once thought to be.

The correct response to this misgiving is to challenge the understanding of these concepts as fossilized lists of approved processes. A society's concept of polite conduct at a time might be given in a list, perhaps.¹⁴ But, the concept of *being objectively justified in believing* (for example) is not some strangely fossilized matter of cognitive etiquette.

It is crucial to keep in mind the standard distinction between concepts and mere conceptualizations. It is also useful to find a few heuristic comparisons. We suggest that evaluative epistemic concepts are much more like the concept of *being nutritious* than like a concept of some fossilized etiquette—for example, the concept of *being ladylike*. There have been *many conceptions* of what makes for being nutritious, many understandings of what nutrients there are and what food possess these in substantial quantities, but these understandings are not competing concepts. Adopting one understanding or conception over another that one previously held is *not a matter of conceptual change*. It is merely a matter of change in one's conception, not a change in concept. One continues to think about the same subject—what is nutritious. One can discover that a certain plant matter works in a certain way in the body—with good developmental effects—and thereby recognize a new nutrient, and thus that the source is nutritious. One is not then changing concepts, but empirically learning something about what is nutritious.¹⁵ Similarly, there have been and remain a diversity of conceptions of what makes for being objectively justified in believing, but these do not represent different concepts. One can discover that certain cognitive processes, under suitable modulational control, make for significant levels of certain forms of reliability and safety. This can amount to a discovery that those processes are objectively justificatory.

Why care? One cares because one *is* an engaged epistemologist—systematically or not, articulately or not, one is involved in suitably modulating one's own truth seeking processes, and one is engaged with others in a community of inquirers. Various evaluative epistemic concepts have grown up around general epistemic values—different constellations of reliability, different constellations of traits useful in the pursuit of systematic true belief. And we today continue to care about these (and other) constellations of epistemically valuable traits.

To an engaged agent, questioning the normative force of an engineering result, or of an application of a familiar evaluative concept, is likely to seem a nonpressing question.

Why care? It is a funny question. It would be rather like asking the engineer designing a reentry vehicle what is the normative force of *being a good insulator* or of the finding that such-and-such a material is a good insulator. It is something obviously valuable in the surface of a vehicle that will need to do what this vehicle is to do. It is like asking a baseball manager what is the normative force of *being a good batter*, or the finding that so-and-so is hard to strike out. Again, for someone engaged with the game of baseball, these traits are obviously valuable. In both cases, the value of in question is pretty obvious, given that one cares about the more general project—winning at baseball or constructing a reentry vehicle. A more challenging question might be why care about these activities. However, the normative practice (of engineering, or of baseball playing or managing, can go forward without articulating an answer to the more abstract question. Engaged agents appreciate the obvious value of the traits.

We have argued that in order to be objectively justified in believing one must form and maintain one's belief by way of processes that are transglobally reliable under suitable modulational control. Being the result of a globally reliable process is not necessary for one's belief being objectively epistemically justified, but something on the order of a concern for global and local reliability is clearly implicated in the kind of ongoing modulational control that is demanded for objective justification. Transglobally reliable processes that would (in a not unusually inhospitable global environment) make for sensitivity to global reliability, and thus for global (and transglobal) reliability under suitable modulational control, is the epistemically desirable feature pivotal to being objectively justified in one's beliefs. This is a feature about which a philosopher might raise the question: "Why care?" To an engaged epistemologist—to one actually engaged with the process of fixing beliefs—the answer should be as obvious as the answer to the question of "why care about on-base percentages" would be to a baseball manager. If there is a question in the vicinity that calls for an answer which turns out to be difficult to articulate, it is the question of why care about forming true beliefs. We will return to such philosopher's questions soon, but for the moment, we want to remain at the more mundane level.

Understanding and employing the common evaluative epistemic concepts may be very useful to any engaged epistemologist—just as understanding the various concepts in play in a branch of engineering may be very helpful in communicating with other engineers and in orienting some of one's engineering practice. One then can better understand and communicate just what is good about a system, and what needs work. Baseball managers, it seems, have

been less prone to confusion about their evaluative concepts. But, to clear headed baseball managers or epistemologists, the answer to “why care” is close enough to being obvious that articulating the answer at the detail philosophers commonly demand seems rather a luxury—the real trick is finding good batters and objectively justificatory processes.

Contemporary philosophers are correct to be engaged with engineering questions focused on local, global, and transglobal reliability, on parallel forms of epistemic safety, and on modulational control contributing to each.

To say that an agent ought to form beliefs in a certain way—to adopt a normative epistemological standard—is to suppose (a) that such agents (typically normal fellow humans) can with supporting training and high motivation form beliefs in the way indicated, and (b) that such processes would represent an optimum or at least satisfactory epistemic way to fix beliefs. To be an optimum or satisfactory epistemic way is a matter of having a fitting tendency with respect to the constitutive end: producing systematic true beliefs. This itself might be understood in terms of the various desiderata we have discussed already: local, global, and transglobal reliability and associated safety, under suitable modulational control. Thus, normative epistemological standards suppose that the processes called for are both tractable and have the general virtues indicated. At least this is so for a kind of all-things-considered normative epistemological system.

There can be normative models attentive to some subset of these general virtues—for example, relative to one form of reliability. Again, we have argued that the concept of *being objectively justified in believing* is centered on transglobal reliability under suitable modulational control. We have argued that, in turn, the suitable modulational control requires an attentiveness to global and local reliability—but that having achieved either is not strictly required of objectively justificatory processes. It requires control processes that are directed to local and global reliability, and that would make changes in other processes as information is accumulated, but does not require that they have as yet achieved modulation enough to have achieved global or local reliability.

Suitable ongoing modulational control—of which naturalized epistemology is a pervasive element—thus requires attentiveness to global and local reliability. The naturalized epistemologist is rightly concerned with all these general virtues.

As also just noted, normative epistemological systems should demand of agents only tractable processes—ones that can be consistently implemented with complementary training and motivation. In the past, philosophers have failed to take sufficient stock of this matter—and their epistemological systems

have commonly been flawed in that they have demanded kinds of processes that ultimately are not tractable for human agents (see Henderson and Horgan 2000). Happily, the temptation to lose sight of tractability of processes is much minimized in the context of naturalized epistemology. The real situated modulational control envisioned here, with its engaged concreteness, does not readily lose sight of what processes are tractable.

In our earlier general discussion of engineering practice and concerns, we noted that engineering work on a problem at a given level will typically be strongly informed and constrained by engineering work and results at other levels. Work on general resources is informed by possible concrete applications of materials or systems with general features. Work on a system with a particular purpose and situated application will draw on general understandings of available materials and possible general system designs and principles. In some cases, local reliability in a specific environment will loom large for the engineer. In others, more robust forms of reliability will be pressing. Situated modulational control of belief fixing processes reflects no less richness of constraints.

The philosopher's questions here are not silly and can be given answers.

All philosophers, like people generally, are engaged epistemologists—that is, to some degree, they are always already engaged in the kinds of situated modulational control that have been discussed here. Some—epistemologists—take this engagement as a matter of central professional concern. Many epistemologists—sharing the philosophical propensity for abstract generalities—would seek to articulate an answer to what might be called philosopher's questions—questions such as, “why should one care about being objectively epistemically justified?” or “why should one care about having knowledge rather than true belief?” Of course, philosophers should not be concerned solely with such questions—rather philosophers should have the diversity of related concerns and questions already found to be a part of engaged epistemology. Still, pursuing these philosopher's questions can serve as a useful way of initiating an inquiry seeking to elucidate the general character of epistemological engagement and value.

As we said, to the engaged engineer the question of why care about insulation on a vehicle that will reenter the earth's atmosphere at high speeds will have an obvious answer. To the baseball manager, the question of why one should care about on-base percentage will seem almost too easy. To an engaged epistemic agent, the question of why care about various forms of epistemic reliability, epistemic safety they afford, and tractability will have similarly easy answers. When one wonders over the “normative force” of

epistemological standards, or of naturalized epistemology, these cannot be the questions that trouble. As we also noted earlier, the question whose answer it difficult to articulate is rather that of why one should care about the project itself: Why care about such vehicles? Why care about winning at baseball? Why care about the production and maintenance of systems of true belief?

If the answer to such questions requires showing that there is some deep value property or properties—a form of good with a “to-be-pursued-ness”—objectively in the world, we despair of answering these questions. (Indeed, we believe that they have no true answer.) This said, it seems to us that there is an answer of a sort to be given to the question of why one should care about producing true systems of belief in a fashion that is locally, globally, and transglobally reliable under suitable modulational control. The answer that we have in mind does not presume to show that systematic true belief is valuable in some metaphysically deep sense. Rather, it seeks simply to locate these virtues as obviously valuable, *given certain presumed value that all readers—indeed, all agents—can be safely presumed to share*. One should care about these epistemic virtues—in effect these engineering virtues of cognitive processes—because one does care about what all agents care about: having true beliefs.

One may wonder whether all agents value having true beliefs—as one will occasionally read a philosopher who denies valuing true beliefs (for example, Stich 1990). However, we believe that such philosophers (insofar as they are agents possessing of beliefs) are just confused. It is a conceptually-grounded necessary truth that all agents value having true beliefs.

Let us reflect for a bit on the concept of *belief*. Without attempting an elaborate analysis here, we can safely say that beliefs are representational states—either occurrent representational states or standing states of a system related to such occurrent states. A representational state (such as belief) is about something in the world (or in some possible world). It represents that thing as being some way, having some feature or property. This is not a wonderfully revealing formulation, but it will do for our purposes here. Our point is not to get much from this remark about representational states generally. Instead, we want to reflect on the point that beliefs are but one sort of representational state, and to note how they are different from other representational states. Our discussion will turn on what distinguishes beliefs from other representational states.

One might say that representational states characterize sets of possible worlds—and that different kinds of representational states turn on taking different stances towards—or adopting different commitments regarding—what is there characterized. Consider some representational states other than belief. One entertains a representation—employs a representation—when *supposing something for the sake of argument*. For example one might grant for the sake of argument that there are “meaning entities,” (or “possible worlds” or “evil do-

ers”) without being committed to there being any. One employs a representation when reflecting on what would result were the world a certain way. For example, one might imagine that the large icecaps covering earth’s land masses melt, and consider what would be the effect on the earth’s oceans. In such cases, one employs representational states that are not beliefs. Belief turns on commitments towards what is represented that are different from the commitments one adopts when merely entertaining something for the sake of argument, or for the sake of considering what would result. The commitment constitutive of belief is such that, if what is represented is not true, one has done something wrong in believing as one does, one has “got it wrong.” This commitment to the truth of what is represented is a hallmark of belief, and distinguishes it from more guarded representation-involving states. Believers, *qua* believers, aim at truth in the sense that they seek to have belief-commitments only with respect to true representation.¹⁶

So, beliefs are not the only kind of representational state—when one imagines some possible state of affairs—say a world championship for the Chicago Cubs—one represents the Cubs as winning. But, one does not thereby believe that the Cubs win the World Series. One might entertain this representation for entertainment, or to project how Cubs fans would react were it to obtain. But, while this may lead one to form beliefs about what would probably happen were the Cubs to “win it all,” (in entertaining these various representations) one does not thereby believe that they have so succeeded, or that the represented consequences have obtained. Roughly, the “purpose” of belief has to do with more than representing a world as being some possible way. All representational states do that much—wishful daydreaming does that no less than belief, or any of the other representational states. The “point and purpose” of belief is to represent what is the case as what is—to represent truths, to hold an is-commitment to a representation that- p only if p . In contrast, one does not sense that one has necessarily done anything wrong when “supposing some proposition for the sake of argument,” even when that proposition is false. The result may yet be instructive in the relevant ways. The point and purpose of daydreaming is to entertain representations that are commonly admittedly false—but whose truth would be in some way gratifying. One gets a certain pleasure from entertaining the propositions, even though some of them may be known to be false—and that is at least part of the point. The daydreaming representations involve a kind of pretend play. (Were some of the propositions not thought to be untrue, or at least thought quite possibly untrue, one could simply take pleasure without the pretense.)

The “purpose” of believing is reflected in, and constituted by, the characteristic “stance” taken by the agent/system regarding what is there represented. By this we mean that the agent, in holding the (descriptive) belief that p , is thereby committed in certain ways—depending on the content of the belief.

The relevant “stance-taking,” or “commitments,” need not be articulated by the agent, nor need they be represented in an occurrent (though perhaps unconscious, or implicit) way. The commitment might be largely dispositional. The dispositions will be exhibited in the felt need to accommodate obviously relevant information—supporting or undermining the ongoing commitment.

At least for creatures like us, the relevant (descriptive-belief-constituting) stance paradigmatically involves *some* openness to the possibility of being corrected—to the possibility of turning out to be mistaken. The openness to being mistaken that is called for here varies with the epistemic character of the content of the representation or proposition in question, and with the character of one’s associated supporting states. That there is such an openness to correction regarding the belief at hand is no less true when the featured representation has to do with a putatively necessary truth such as a mathematical truth or a conceptually grounded necessary truth. One can get it wrong in connection such matters. Frege was fittingly open to being mistaken regarding to the principle of class abstraction—although it had seemed so obvious to him that he was once ready to depend on it as a central piece of his system. Thus, his response to Russell’s paradoxes. Similarly, even when one has what one initially judges to be a good proof for a mathematical belief, the reservations of the experts in the mathematical community may give one pause, and induce one to look more carefully at the proof (Kitcher 1983). The openness of the scientific community to theory revision stands as a paradigmatic form of openness to correction involving empirical propositions. In all cases, being utterly closed to the possibility of being mistaken, being utterly closed off from rethinking matters even in the face of conflicting evidence,¹⁷ suggests that one is not aiming at “getting it right” in the way descriptive beliefs characteristically must—that one is not aiming at the truth in one’s representational states.

By way of contrast, consider some cases in which one would not be open to being mistaken and one is not directly aiming at truth. Perhaps one is imagining the scenario in question, or arbitrarily supposing it for the purpose of generating a formal system with a given formal property. In making the supposition itself, one is not aiming to truthfully represent the world (although one indirectly aim to discover truths about a certain class of formal systems in which the supposition would be satisfied). Think of the early investigations of non-Euclidean geometries—where one might suppose that for any line on a plane, and any point off that line, there were multiple lines (or no line) parallel to the initial line and passing through the point. One might do this without prejudice or commitment to whether the geometrical structure of the world was as there supposed. One’s aim is something on the order of conditional proof, not to describe the world in that supposition itself.

It is intrinsic to descriptive belief-states that one is committed to getting it right, and thus somewhat sensitive to the possibility of getting it wrong, mis-

representing the world, on matters there represented. This much alone—the contrast between descriptive belief and the more guarded commitments associated with other representation-involving states—seems to indicate that at least clear cases of belief are associated with a concern for, or “aiming at,” truth—getting it right, correctly representing the world. One must have this concern as a part of one’s stance toward the representation—and this concern must at least involve (perhaps be constituted by at least) dispositions to reconsider, dispositions that involve some content appropriate readiness to reconsider the claim in questions or related claims. Thus, to have a belief is to aim at truth, to be concerned with the truth of the relevant representation—at least in an integral, if perhaps inarticulate and background, way. This concern is partially constitutive of believing.

One can approach these issues in a slightly different way. This also seems to us a conceptually grounded necessary truth: belief states interact with each other, and with desiring/valuing states in certain characteristic ways: Ways that crudely approximate practical syllogisms, for example. Ways that approximately conform to what are “obvious” inferential patterns—practical or “theoretical.” The pattern of the interaction of belief states with each other, and with desiring states, is a matter of conforming at least to “obvious” inferential patterns. What is “obvious” must be understood so as to make due allowance for common errors made in connection with more sophisticated normative standards. What is “obvious” in the relevant fashion is not an *a priori* matter.¹⁸ Let us focus on certain inferential patterns involving talk of belief associated with “theoretical” reasoning—reasoning about what is the case. Among the most obvious patterns are those that might be thought of as the avoidance of beliefs that would constitute an obvious instance of Moore’s paradox: p , but I don’t believe that p . Include here the variant on Moore’s paradox: $\sim p$, but I believe To earnestly think, “ $\sim p$, but I believe p ,” immediately would prompt a revisionary reflection on just what it is that one really believed—apparently not p . To encounter an agent earnestly insisting on something amounting to an instance of these paradoxical formulations would lead one to wonder what it is that that agent really believes. There is no parallel avoidance to be found in connection to supposing for the sake of argument, in which one considers what would be the case were p the case. Without discomfort or embarrassment, one might think: “To see what would follow, suppose that p , although I do not believe p ” The avoidance of Moorean-paradoxical beliefs seems to reflect the way in which belief—full-fledged descriptive belief—aims at truth. The paradoxical formulations express conjunctions. Consider the formulation: “ p , but I don’t believe p ” Of course, such formulations might merely express in a picturesque way a dialectically guarded supposition—something less than belief. But, suppose that it is to express a belief. What then is this belief? The first conjunct is simply p . So, the conjunctive belief

would involve the belief that p , and more. The second conjunct is then the denial that one believes p . This conjunct is doubly puzzling. Perhaps one can understand how someone might believe something, and not know it. Such is certainly not a paradigm case of full belief—but it is perhaps understandable as a degenerate case. But, it is much more difficult to understand how an agent might believe something, do so in a way that allows that agent to sincerely assert that which is believed, and yet be so oblivious to this belief that that agent would in the same breath deny holding the belief. One simply cannot understand this scenario—it does not seem to involve a state that interacts with others as beliefs do. Now consider the formulation: $\sim p$, but I believe p . It is, if anything, more problematic. The first conjunct seems to express the belief that $\sim p$. The second conjunct expresses the belief that one believes that p . Again, it is perhaps understandable that one could hold both beliefs in some fashion in which they are insulated from each other. But, it seems mysterious how one could believe both in the same space of reasons—in the same cognitive “breath.” This is reflected in the instance of belief revision, when one notes that one believes p and then goes on to realize that $\sim p$, one at that instant has noted that one “had it wrong,” and thereby revises one’s previous belief. The point of belief, is to “get it right,” to represent the truth about how things are. (See also [Railton 1996](#), 131 and 133.)

We have argued that the point and purpose of belief is to represent the truth (at least for non-normative beliefs). This is to say that one who believes has thereby the desire to “get it right”—to get at the truth of the matter. All believers—*qua* believers—value truth. It is then a short step from this to the result that all agents value truth—for agents must, it seems, have beliefs. Nothing could be an agent that did not possess beliefs and desires. This provides an answer to “the philosopher’s question” (“Why care?”) that situates naturalized epistemology and epistemology generally—allowing them to be understood as normative cognitive engineering, as engaged modulational control in the interest of producing true systems of belief. The answer is that one cares about such modulational control, such cognitive engineering, such self-regulation and regulation of one’s epistemic community, because one cares about the project of arriving at systems of true belief.

Notes

¹ Bishop and Trout (2005b, 696) say that, “a naturalistic theory of epistemology takes as its core, as its starting point, an empirical theory.” Notably, the empirical work or theory to which they themselves commonly appeal, what they term Ameliorative Psychology, is itself unabashedly normative. We would prefer to see this work as itself a piece of ongoing naturalized epistemology. When one asks just what really is the “starting points” here, one

is likely to find them somewhat diverse. Honestly, the notion of a starting point seems somewhat obscure here.

- ² Note well: A possible global environment can be compatible with one's having such experiences within it even if these experiences are radically and systematically nonveridical.
- ³ Given your lack of information and training, there were no epistemically desirable processes open to you here. (Well perhaps one: one might perhaps receive testimony from better trained colleagues over field radio.) But, while you may need to do something, you need not form a belief—and we are stipulating that you do so here by guessing.
- ⁴ The contemporary interest in the epistemology of testimony reflects the importance of testimony in our epistemic lives. Accordingly, one could socialize the above point: optimum human epistemic performance is likely optimum community performance, where what are yoked and pulling in tandem are not just a given individual's cognitive processes, but also the various complementary individuals comprising the epistemic community.
- ⁵ Alternatively, one insists that these conclusions be supported by double-blind studies with controls employing otherwise understood treatment alternatives.
- ⁶ A belief-forming process P may be under the conditioning control of a wider set of processes, with or without those wider processes having yet come by information that prompts changes in, or modulations of, P. When there is such a functional relationship between processes we will say that the process P is *under the modulatory control* of the wider processes. This wider set of processes may be termed *conditioning processes* with respect to P. Being under the modulational control of wider processes and being a conditioning process is a matter of the dispositional or control relationship between wider processes and some narrower processes.

When the conditioning processes then turn up information, or apparent information, bearing on the reliability of a process P, and when P or its use thus comes to be spawned, tailored, selectively triggered or inhibited, or in some like manner refined, we will say that P is *modulated* by those conditioning processes, S.

- ⁷ Issues of tractability provide one way of thinking about the statistical prediction rules that so interest Bishop and Trout (2005a). Consider Goldberg's Rule. Mental health professionals might try to integrate diverse information in a way that modulates their categorization processes for neurotics and psychotics. At least given the character of available information, the research supporting Goldberg's Rule, seems to indicate that this task is relatively intractable for folk, and that the SPR provides the more tractable and thus more robustly reliable process. Psychological investigations of anchoring in judgments of probability, for example, also suggest that certain forms of modulation will be difficult for human agents.
- ⁸ For concreteness, here is a possible case. Perhaps the agent, Luke, lives in great proximity to the Smoky Mountain National Park and adjacent Cherokee National forest. Luke learned to identify many trees in those forests. He has a process for identifying the White Pine (*Pinus strobes*): count the needles in a representative set of bundles. If bundled in groups of five, the pine is judged a white pine. Because the white pine is the only pine in those forests with five needles to a bundle, this process is locally reliable. But, were the process applied in other local environments it would be highly unreliable. In the mountains of Wyoming, for example, both the limber pine (*Pinus flexilis*) and the whitebark pine (*Pinus albicaulis*) are numerous and have five needles to a bundle. If Luke has only learned to associate five needled bundles with the white pine, and nothing more, say about distribution, his process is merely locally reliable.

- ⁹ Suppose for example that Luke's brother, George, began with training and a locally reliable categorization process like Luke's. But, hearing that tree species have ranges, George acquires a set of field guides as he travels. These guides are the product of careful comparative observation by professional botanists and naturalists, and are produced by processes that are globally (and transglobally) reliable. His five-needle rule for white pines will be inhibited when not in local environments such as the Smoky Mountains. Complementary rules or process will be developed for other regions. Under such modulation, the refined process of white pine identification becomes globally as well as locally reliable.
- ¹⁰ Bishop and Trout (2005b, 696) say, "a naturalistic theory of epistemology takes as its core, as its starting point, an empirical theory." Notably, the empirical work or theory to which they themselves commonly appeal, what they term Ameliorative Psychology, is itself unabashedly normative. We would prefer to see this work as itself a piece of ongoing naturalized epistemology. When one asks just what really is the "starting points" here, one is likely to find them somewhat diverse. Honestly, the notion of a starting point seems somewhat obscure here.
- ¹¹ This guarded formulation reflects the authors' rather revisionary understanding of much reflection as having itself an ineliminable empirical dimension that leads us to write of a kind of "low-grade a priori" (Henderson and Horgan, SJP, opulent)
- ¹² At least, for categorizing those who can meaningfully take the MMPI. Of course, this claim of global rather than local reliability itself turns on judgment regarding how widely Goldberg's results generalize. To apply the rule outside populations of western patients, should be supported, it seems, by empirical research—otherwise one's process seems unacceptably modulated.
- ¹³ In cognitive agents such as humans, some significant range of modulational control processes work by way of an articulate concern for global and local reliability. To the extent that this holds, one can say that a *concern* for such forms of reliability is epistemically required as an epistemic means to the attainment of the constitutive epistemic goal. But, it also seems that some modulational control is managed by way of a sensitivity to the local and global reliability that never rises to the level of articulate concern. Thus, the concern of which we here write should be understood broadly. Alternatively, one could say that a kind of attentiveness or sensitivity to local and global reliability is required for an agent to qualify as having suitable modulational control processes in place.
- ¹⁴ There are reasons for thinking that even this would be too flat-footed an understanding of etiquette.
- ¹⁵ One senses that differing conceptions of gendered etiquette that might reasonably be taken as characterizing what it is to be "ladylike" are different. Suppose that the Victorian conception of being ladylike has some crude parallel in the etiquette of the Yanomamö—a conception that effectively constitutes what makes for proper social self presentation for a woman in the society. It would seem silly to suggest that the one conception or the other is correct—and that one society had discovered something about what it was to be ladylike that the other had missed. One senses that there is no shared subject and no shared concept. If a woman were to "go native," moving from the one society to the other, one would be inclined to say that here she changed concepts rather than mere conceptions.
- ¹⁶ However, we should note at the onset a qualification. [Horgan and Timmons \(2000\)](#) argue that moral beliefs—evaluative beliefs generally—work differently from descriptive beliefs. Both kinds of belief involve commitments with respect to "a core descriptive content," (p. 132) but rather different commitments. Consider the descriptive belief that terrorists flew airliners into a large office building, and closely related evaluative or moral belief that it

was wrong for terrorists to fly airliners into a large office building. These beliefs share the same core descriptive content, expressible by the that-clause, *that terrorists fly airliners into a large public office building*. Both employ a representation with such cognitive content. However, the descriptive belief involves an “is-commitment” with respect to that descriptive content, while the evaluative belief involves an “ought-commitment” (or, rather, “ought-not-commitment”) with respect to that content. In the case of a moral belief, one has done nothing wrong, when the that-clause (the representation) is false. Clearly, in the case of the descriptive belief, one has got it wrong when the that-clause is false. Despite these differences in commitments regarding what is represented, Horgan and Timmons (2000) argue that both are beliefs. They propose an understanding of ought-commitments according to which these do not aim at the truth of any representation—although, by virtue of their stance with respect to the descriptive content (the that-clause), ought commitments do have cognitive content. If this “nondescriptive cognitivism” is correct, then the remarks of the present discussion regarding beliefs having truth as an intrinsic aim would need to be understood as treating of descriptive beliefs with their is-commitments. We will so understand them.

- ¹⁷ There are interesting issues one might pursue regarding dogmatic belief—perhaps of a religious sort. We cannot pursue these here except to say this much. Even when a belief is held dogmatically, the holder commonly has response tendencies involving the defensive belittling of apparently undermining evidence.
- ¹⁸ Again, we come to issues having to do with the character and limits of what has been termed “charity” in interpretation. See also, [Quine \(1986\)](#) and [Henderson \(1988, 1993\)](#).