

Epistemic Unification

Mitchell R. Haney

Missouri Western State College

Herman E. Stark

South Suburban College

Much epistemological theorizing is the attempt to specify what makes for meritorious cognition, but epistemologists have not, despite meritorious effort, achieved unity when it comes to picking out the feature and principles that are distinctive of epistemic normativity. In this essay we explain why this is the inevitable outcome. We isolate important but overlooked variations in the link between epistemological theorizing and the idea of epistemic unification, and then argue that much epistemological theorizing is misguided because it aims toward complete epistemic unification when only partial epistemic unification, at best, is possible. But our arguments — based on work in moral epistemology and philosophical psychology — stop short of epistemological eliminativism, and thus we stake out a middle ground between philosophers such as Descartes, the earlier Alston, Audi, and the earlier BonJour on one hand and Rorty, Fish, and Patricia Churchland on the other.

Much epistemological theorizing has been and is the attempt to advance *content* for both “the feature” and “the principles” distinctive of epistemic normativity. Epistemologists propose various features and principles, show how they accommodate certain *prima facie* cases of epistemically meritorious cognition, and then brace for the inevitable counter-examples. Consider the argot: self-evident, foundational, coherent, reliably-produced, conditions of justification, and principles of rationality. Apart from this, however, there is

We thank David Waller, Hal Brown, Ruth Marcus, Ray Russ, anonymous reviewers, and members of APA and World Congress audiences for their corrections on earlier versions of this essay. Requests for reprints should be sent to Mitchell R. Haney, Ph.D., Department of Humanities, Missouri Western State College, 4525 Downs Drive, St. Joseph, Missouri 64507 or Herman E. Stark, Ph.D., Department of Humanities, South Suburban College, 15800 S. State Street, South Holland, Illinois 60473.

little that is conclusive to report concerning this type of epistemology — except, of course, distinctions, disagreement, and discontent.

In recent years the topic of epistemic unification has been addressed in philosophical presentations, publications, and even symposia.¹ Not surprisingly, it turns out that epistemologists are not unified in what they mean by epistemic unification, whether it can be achieved, and even why it matters. In this paper we nonetheless jump into the fray because we see something in the topic of epistemic unification that matters profoundly to epistemological theorizing, and indeed something that has not been adequately isolated, much less underlined, despite the current interest in epistemic unification.

We understand epistemic unification to be one major goal of epistemological theorizing. Insofar as epistemological theorizing attempts to produce a general account for the full range of “meritorious thinking,” that is, epistemic normativity, it is aiming to achieve epistemic unification. The something about epistemic unification that has not been adequately isolated, much less underlined, and that matters profoundly, is the cluster of variations in the link between epistemic unification and epistemological theorizing. Variations have been missed, we propose, because one strand has been confused for the link. The usurping strand is that of “ideal connection,” that is, epistemological theorizing that intimates, or is guided by, the idea that epistemic normativity takes the form of a structure that is specifiable as a maximally simple, maximally complete system. Epistemological theorizing, in short, tends to aim at “ideally-complete epistemic unification.”

The idea of such an exhaustive structure captures the extreme sense of complete epistemic unification, and so it is useful to distinguish it from both “non-ideally-complete” and “partially-complete” structures. Minor deviations such as “relatively simple” or “almost complete” would make for structures that are non-ideally complete and partially-complete, respectively. Moreover, it seems to us that the optimal, if not indeed the only, way to express such a maximally simple, maximally complete structure is in terms of a single, simple, and general epistemic feature and a compact set of epistemic principles. And the possibility of this, that is, of the consequent of the only-if relation just hinted at, is what we deny.

Our purpose in this paper is to set forth and defend a meta-level critique of the epistemological theorizing that aims at ideally-complete epistemic unification. Our main conclusion is that numerous theories of epistemic justification and rationality were and are bound to be inadequate because they are built on the unfulfillable hope that an ideally-complete account of epistemic normativity is possible. *We deny that ideally-complete epistemic unification is*

¹Note, for example, the very title of the 1996 SUNY Buffalo Marvin Farber Conference: Can Epistemology Be Unified?

possible, that is, that the full range of epistemically meritorious cognition is expressible solely in terms of a single, simple and general epistemic feature and a compact set of exceptionless epistemic principles; thus, we also deny that epistemic normativity can be completely unified in a maximally simple, maximally complete (non-disjunctive) structure. We hold, accordingly, that “the structure of epistemic normativity” is decisive when it comes to the issue of epistemological unification and the possibilities for epistemological theorizing. But we are not epistemological eliminativists either. We contend that a sort of epistemic unification is possible, and that a version of “feature and principles” epistemological theorizing is genuinely pursuable. We end by proposing, in short, a legitimate, albeit lower, aim for this kind of epistemological theorizing, but with the caveat that we do not necessarily endorse this kind over other conceptions of epistemological theorizing, and thus end up staking out a middle ground position between philosophers such as Descartes, the earlier Alston, Audi, and the earlier BonJour on one hand and Rorty, Fish, and Patricia Churchland on the other.

Epistemological Theorizing

One central aim of epistemology is to provide an analysis of *epistemic normativity*. More precisely, epistemologists seek to provide the correct theory — a *compact* but *general* theory — as to what constitutes epistemically normative cognition. The familiar controversies over the conditions of epistemic justification provide the major illustration of this aim, but another is provided by the more recent discussions on the nature of rationality. We note, however, that in the various proposed theories one typically finds that epistemologists attempt to explicate epistemic normativity, though not always by this name, by identifying a simple and general epistemic feature and enumerating a compact set of exceptionless epistemic principles. This is true of foundationalist, coherentist, and reliabilist works, to name just some mainstays from among the numerous current epistemological theories, and a brief meta-level look at some representative examples will show that what they tend to have in *common* is more significant than what they do not.

To begin with foundationalism, Robert Audi (1993, p. 134) has developed a “fallibilist foundationalist” version that would anchor epistemic normativity in cognition that involves beliefs for which there is no reasonable doubt. Thus, the *feature* distinctive of epistemic normativity would be something like “involves a foundational belief.” And since logical connectedness and coherence play roles in determining whether a belief is foundational or adequately connected to a foundational belief, the epistemic principles bound up with epistemic normativity would involve the rules of implication and the idea that incoherence is an underminer (p. 135).

We should note that we use “distinctive” throughout this paper so as to avoid disagreements as to whether such features are necessary, sufficient, or necessary *and* sufficient for epistemic normativity. Many epistemologists would claim that their feature is necessary and sufficient for epistemic normativity, but Audi’s use of incoherence as a feature that is sufficient for undermining a belief, even if the belief is adequately grounded to a foundational belief, is an example of treating the foundational feature as merely necessary.

Next, Laurence Bonjour’s (1985) well-known defense of a coherentist structure of justification for empirical knowledge can also be seen, on the meta-level, as an attempt to fill in the blanks for the epistemic feature and the epistemic principles. Bonjour’s candidate for the feature distinctive of epistemic normativity is coherence, where coherence would somehow involve not only logical consistency between beliefs but also a consistency between beliefs and empirical inputs (p. 118). The epistemic principles thus would include “One ought to seek logically consistent belief sets” and “One ought to trust sense-inputs under normal conditions” (p. 118). And as for the facts that Bonjour’s account seems restricted to empirical beliefs and that he more recently seems to have backed off of coherentism (1997), these only feed into our overall thesis that no single, simple, and general feature can capture the full range of epistemically normative cognition (or even, as we argue below, all levels of epistemic normativity within a single domain).

Lastly, Alvin Goldman’s (1979, p. 2) reliabilism is based on the idea that “being produced by a cognitively reliable process” is the feature distinctive of epistemic normativity. His candidates for the epistemic principles would thus include “Under normal conditions of light, one ought to trust one’s visual inputs” and “Under normal mental conditions, one ought to trust one’s reasoning” (p. 19).

The overall point, then, is that a meta-level look at some major epistemological theories reveals that they intimate — if not are guided by — a common ideal. The common ideal is that of explicating epistemic normativity by identifying a simple and general epistemic feature and a compact set of epistemic principles. What some major epistemological theories *share*, in other words, is a view of what the form — sans content — of epistemic normativity is. As for why this ideal is held, the reasons are clear enough: if such an account is achieved, for example, the actual (*pace* Bonjour, 1985) and potential instances of epistemically normative cognition would be systematized in a tractable, and indeed manageable, way. But our point is to note that any attempt to achieve tractability and indeed manageability in this way, that is, via a simple and general epistemic feature and a compact set of epistemic principles, trusts that the full range of epistemic normativity fits onto a “rigid” structure, that is, a structure specifiable as a maximally simple,

maximally complete system. We consider below whether there might be other ways of achieving tractability or manageability that avoid the rigid structure, but the present concern is to underscore that “competing theories” are not necessarily or even actually competing on structural matters. The differences between the theories, in other words, arise when it comes to pushing candidates for the content of the simple and general feature and the compact set of principles that are thought to capture epistemic normativity; theories may differ in terms of how they “fill in the blanks” without differing with respect to their commitment to the generic *form* underneath.

The Two Assumptions

The foregoing remarks about epistemological theorizing pointed to the structure of epistemic normativity, and now we will formulate more precisely what much epistemological theorizing either intimates or presupposes about the structure of epistemic normativity. We claim that much epistemological theorizing rests, whether implicitly or explicitly, on the following two assumptions about the structure of epistemic normativity.

The complete systematicity assumption. Epistemic normativity is completely systematizable by a simple and general epistemic feature. This means that there is some non-hybrid feature, for example, “the result of a cognitively reliable process,” that is in some tractable, manageable and systematizing way distinctive of all epistemically normative cognition.

The complete codifiability assumption. Epistemic normativity is completely codifiable by a compact set of exceptionless epistemic principles that are sufficiently non-vacuous to serve as cognitive directives or illuminating explications. This means that there is some manageable group of exceptionless general statements that usefully describe what makes for epistemic normativity, for example, “Beliefs that follow from observation under standard conditions are beliefs that are the result of a cognitively reliable process.” The relative clause immediately preceding the example serves to indicate what an epistemic principle is supposed to do, and also to rule out truisms such as “Maximize true beliefs and minimize false beliefs” or even “One should think carefully.” For purposes of contrast we suggest the following as an example of a *non-exceptionless* epistemic principle: “Beliefs that follow from observation under standard conditions are, *ceteris paribus*, beliefs that are the result of a cognitively reliable process.”

A classic example of these assumptions working together is found in Descartes’ attempt to provide “rules for the direction of the mind.” According to Descartes, the feature distinctive of epistemic normativity is self-certification, and the epistemic principles direct cognition toward that which is clear and distinct and — according to a popular reading of Descartes — toward

that which is deducible therefrom. A better reading is to see self-certification as the feature against which something is *compared* to, and not deduced from, in order to determine the presence of epistemic normativity, but for present purposes the popular reading furnishes a cleaner illustration. Descartes assumes that epistemic normativity takes the form of a completely unifiable structure in that he attempts to account for the range of epistemic normativity in terms of a feature, self-certification, and principles that involve, for example, the rules of deduction.

It should not be difficult to see that the complete systematicity and complete codifiability assumptions jointly accommodate the possibility of achieving ideally-complete epistemological unification. *The epistemologist could achieve such unification by correctly identifying the systematizing feature and the set of principles that are distinctive of epistemic normativity.* It is a more telling discovery, however, to note that the falsity of either assumption implies the impossibility of full-blown epistemological unification. This is more telling because few epistemologists note explicitly the assumptions, much less what follows from their falsity. But be that as it may, our contention is that if either assumption is false, then a major implicit reason for thinking that complete epistemological unification is achievable has been undercut. Even stronger, we contend that establishing the falsity of either assumption establishes the impossibility of ideally-complete epistemological unification. If there is no single, simple, and general epistemic feature, or no compact set of exceptionless epistemic principles, then the correct epistemological theory will contain either non-ideal complexity, irreducible plurality or non-exceptionless codification.

Showing the Falsity of the Assumptions

We reject both the complete systematicity and complete codifiability assumptions, but since we believe that a fair amount of epistemological theorizing is wittingly or unwittingly guided by them, we will offer arguments against these assumptions — *even though we have not found attempts to argue for them.* Our argument against the complete systematicity assumption is analogous to arguments posed by W.D. Ross and “moral particularists” against moral monism as an element of the structure of moral normativity. But our argument differs in that we use “features arguments” against features appropriate of epistemic normativity and not of moral normativity. Our argument against the complete codifiability assumption will be based on the frame problem (a problem discussed in philosophical psychology). We will present our case against the complete codifiability assumption first because our case against the complete systematicity assumption will more naturally lead to our subsequent positive claims concerning both the structure of epistemic normativity and the prospects of epistemic unification.

The Falsity of the Complete Codifiability Assumption

The following argument uses reflections on the nature of cognition as the basis for the claim that any proposed compact set of exceptionless epistemic principles is doomed to be, at best, only slightly better than a collection of epistemic truisms. The central premise is that there is a component process of some, if not all, instances of epistemically normative cognition that is too subtle to be captured by a compact set of exceptionless *processing* rules. The process involves determining relevance, and it is *invariably presupposed* by proposed codifications of epistemic normativity. From this premise — that is, the premise that the relevance-sensitive cognition required for competent belief updating outstrips exceptionless processing rules — we conclude that it is wrongheaded to attempt to formulate a compact set of exceptionless epistemic principles for the purposes of codifying epistemic normativity, for we suggest that what is psychologically impossible limits what is epistemologically feasible. We take the premise to reveal, in other words, that there is significantly less directive or explicative force in proposed epistemic principles than epistemologists tend to realize, and that illustrations of the point can be had by looking “underneath” proposed sets of epistemic principles to draw out the naturalistically naive assumptions about how epistemically normative cognition works.

This component process of epistemically normative cognition that cannot be captured by a compact set of exceptionless processing rules is not a “minor module” but rather a highly central process. The process is that of updating beliefs in light of new information, and the problem of finding a rational method for doing this has become known, after one such attempted method, as *the frame problem*.

The frame problem is much discussed in the cognitive science literature (see, for example, the collection of analyses in Pylyshyn, 1987, and for more recent work, Haselager, 1997). For present purposes we suggest understanding the frame problem as the problem of getting a cognitive system to determine the relevance — if any and how much — of new information to the system’s current beliefs. Many humans manage to do such cognitive feats quite well, and of course without consciously consulting exceptionless processing rules. But the following discussion underscores difficulties of an “in-principle” sort that arise when attempting to solve the frame problem via processing rules, irrespective of how the rules are consulted.

There are a variety of approaches to the frame problem. Pylyshyn, for example, classifies the frame problem as a particular kind of holism problem (1987, pp. viii–ix), and this is not an unfamiliar view of the matter. But Haselager (1997, p. 144) has more recently distinguished between and explored logical and psychological approaches to the frame problem, and

claims that on the psychological approach, it is not necessary — in order to solve the frame problem — to solve problems of inductive logic. We are impressed by Haselager's analysis, but present purposes follow Pylyshyn's lead and situate the frame problem within the more general holism problem, and accordingly first introduce the holism problem before presenting our analysis of the more specific frame problem.

The holism problem arises when one attempts to model general reasoning. General reasoning, as opposed to reasoning within a closed domain, for example, chess, *requires that a cognitive system be able to connect any single representation (belief) with any other.* Now, many humans have this ability, and some like Sherlock Holmes are brilliant at it, but it is also clear that humans do not determine the *relevance* of any one representation to others by making all possible inferences from it or their other representations. This is a problem for attempts to model general reasoning because

... there appears to be no way *in general* to index beliefs so as to exclude certain obviously irrelevant paths of inference in advance, without at the same time excluding some obviously relevant ones. On the other hand, if "irrelevant" inference attempts are not suppressed, the cognitive system will become *mired in an exponentially expanding range of inferences* and will never be able to deal intelligently with domains requiring a potentially unbounded domain of knowledge — *such as is involved in carrying on the simplest everyday conversation* [emphases added]. (Pylyshyn, 1987, p. viii)

The germane point is that from any particular representation there are many chains of inference leading to *countless* consequent representations, some of which are relevant, but there is no manageable way to determine *in advance* which chains lead to these relevant consequences. A chain relevant in one context need not be in another, and thus, to jump ahead, the attempt to capture competent belief updating via a compact set of exceptionless processing rules is fundamentally problematic — especially if such processing involves making relevance connections via inference in "an exponentially expanding range."

The frame problem is a more specific case of this general holism problem. The holism problem is that of relevantly connecting any representation to any other, and the frame problem is that of finding a rational way of determining what to change and what to leave the same *in the face of input*. Since reasoning about what requires updating (given an input) can involve *any* representation in the whole system, it seems that such "coherence-seeking" cognition requires an exponentially explosive number of inferences to be made. But that is not correct, for one can quickly and easily update beliefs upon hearing a dinner companion asking the waiter for salt — and one can do this without having to go through the history of chemistry, of money, or remembering one's arcadian days in Salzburg.

One approach to this belief-updating problem involves developing a formalism that included “frame axioms.” These frame axioms are “statements that specified which properties of the world would remain unchanged when a certain action was carried out” (Pylyshyn, 1987, p. ix). But it became clear that this solution does not avoid the old problems. Consider, for example, all that might change in a daily plan upon seeing the neighborhood pest in the mall. Would it now be a better plan to leave the mall and come back later? Should one take this as the last straw and forgo buying the sweater altogether? Or just for the time being? Should one use this as an excuse to take a cappuccino break? Should one go buy a book on fate? Should one abandon the belief that one is not being followed? Should one jettison the belief that the pest does not play hooky from work? It thus seems implausible that frames, that is, instructions which isolate the beliefs that may require revision in the light of new information, are the answer. It seems implausible, in other words, to believe that one can spell out in terms of a *compact set of exceptionless processing rules* all the potentially relevant factors that rational cognitive systems must and do track.

The above discussion highlights the fact that the cognitive feat of spotting relevance is highly contextual and can involve too many factors for “relevance-sensitive” processing even to be *conformable* to a compact set of exceptionless processing rules, and *a fortiori* for it to be achievable by “unconsciously consulting” such rules. Even further, the point extends to the idea of procedural knowledge, i.e., implicit, non-occurrent information that is capable of being fully represented in “long-hand” (e.g., “trained-in” information), for the problem underscored here has to do with principles, not with the manner of principle consultation, whether that manner be conscious, unconscious, or a hybrid (see Henderson and Horgan, 2000, for further discussion on procedural knowledge). The upshot, then, is a challenge to the complete codifiability assumption, for the impossibility of formulating a compact set of exceptionless processing rules for relevance-sensitive cognition suggests — via the “ought implies can” idea — that we ought not seek to formulate a compact set of exceptionless epistemic principles. We can stop short, by the way, of suggesting that an outright implication holds between the nature of cognitive processes and epistemic norms and instead maintain minimally that something like an analogy holds. Horgan and Tienson seem to take a similar position:

If the cognitive processes that lead to moral and other normative judgments are soft, however, then there is no particular reason why ethical, epistemic, or other normative standards employed by or employable by humans should be systematizable via general, exceptionless principles or rules. (1996, p. 143)

The key claims are now in place, and we can now sketch the argument. First, insofar as epistemic principles are to serve as useful directives or expli-

cations for epistemic normativity, the “relevance-tracking” requirement on competent cognition, as underscored by the frame problem, stands prohibitively in the way. No compact set of epistemic principles can hope to direct or explicate sufficiently the relevance-sensitive part of cognition that is required for both the proper implementation of any single epistemic principle and the adjudications between competing principles of the set. The epistemic principle, “Under normal mental conditions, one ought to trust one’s reasoning,” for example, requires determinations of what is and is not relevant to normal mental conditions, and furthermore the principle may not be sufficient for the epistemic normativity of a result if other epistemic principles are relevant, for example, a counter-result by a community of experts might make an “expert consultation” principle relevant, or a counter-result based on one’s visual observations under standard conditions might make another reliabilist principle relevant. Second, the relevance-tracking requirement on competent cognition is prohibitive to epistemic principles in the sense of directives or explications that are always necessary for epistemic normativity. The kind of example just mentioned can illustrate against not only the alleged sufficiency of the directives but also the alleged necessity of the directives. There is, then, something grossly deceptive about advancing a set of principles for epistemic normativity when the set pastes over the very difficult matter of how cognition is to proceed in order to achieve conformity to the principles; connecting a belief to a self-certifying belief, for example, requires that one locate the *relevant* deductive path from the countless deductive paths leading from any given belief.

As a coda, suppose one grants the above but attempts to use it against our conclusion by suggesting “Determine relevance” be added as an epistemic principle. The proposed principle, ironically, is one that we like to endorse as a fine target for epistemological theorizing, but our main response is that the frame problem shows that the proposed principle — necessarily a general statement, *since generality is needed for compactness* — is too vacuous to serve as an informative directive or explication. Our response, in other words, is to make explicit a lurking trade-off. On one hand, the more vacuous a principle is, whether it be understood as necessary or sufficient for epistemic normativity, the less interest will it be to epistemologists trying to capture epistemic normativity in terms of illuminating principles (for example, algorithmic or heuristic directives for cognition). On the other hand, the more non-vacuous a principle is, the less likely it is to be non-overrideable. The overall point, then, can be formulated in these terms: one cannot completely codify the full range of epistemic normativity via a compact set of exceptionless epistemic principles because a significant component of competent cognition that is presupposed by them cannot be so codified, and thus this argument alone rebuts the goal of ideally-complete epistemic unifi-

cation — even if there were a single and general feature distinctive of epistemic normativity.

The Falsity of the Complete Systematicity Assumption

The following arguments strongly suggest that there is no simple and general feature that is distinctive of all of epistemic normativity. This is to say that in some instances of epistemically normative cognition a feature that is otherwise indicative of epistemic normativity (“Produced by a cognitively reliable process”) is not present — or is not indicative even though it is present, or is present but not indicative to the same degree.

To begin, we note that moral particularists who claim that there is no simple and general feature in virtue of which acts have the moral properties they have are thereby making a claim about the structure of moral normativity. They argue, in effect, that there is no feature that is either always morally relevant or always morally relevant in the same way each and every time it is instantiated. For example, many consequentialists hold that pleasure is always morally relevant and always a good or right-making feature. Particularists are prone to counter that a feature like pleasure is not universally relevant to the morality of actions or character traits, because there are cases where pleasure is instantiated but either is not morally relevant or is morally relevant in a way that is not indicative of moral goodness. The fact that an offender attained pleasure by slaying the victim does not mitigate — even slightly — the wrongness of the action; indeed, the fact that the slayer experienced pleasure as a result of doing the action intensifies our judgment that the slayer is wicked, and thus it is clearly not a good thing at all, in such a case, that the person took pleasure in the action. On the basis of such examples moral particularists are led to conclude that pleasure is not always morally relevant and that it is not always a good or right-making feature. Similar counter-examples against the variety of proposed moral criteria — criteria that are infamous in debates between competing ethical theories — suggest that there is not a simple and general feature in virtue of which acts and character traits have the moral properties that they do. The minimal suggestion that arises from such counter-examples is that there must be a plurality of features that determines moral normativity. A parallel suggestion can be made concerning the features in virtue of which beliefs are epistemically normative.

Much epistemological theorizing presupposes that there is some simple and general feature in virtue of which beliefs — in many realms of discourse and fields of inquiry — are epistemically justified, warranted, or rational. But we maintain that worries raised by moral particularists against generalism in ethics, especially of a monistic variety, apply analogously against all episte-

mological theorizing that presupposes the complete systematicity assumption. If there are intuitively viable counter-examples to all proposed epistemically-distinctive features, then there is good reason to believe that the complete systematicity assumption is an unwarranted assumption and that the goal of complete unification within epistemological theorizing is wrongheaded. But we will here go beyond the healthy supply of extant counter-examples and instead offer two *general* counter-examples that, we think, indicate the falsity of the complete systematicity assumption. This stratagem of producing sweeping counter-examples stands in contrast to the common practice of providing individual counter-examples aimed against specific proposed features. Hilary Putnam (1983, pp. 233–234), for instance, argues that reliability is insufficient for justification, because then any process — including things like Ouija boards, consulting a perfectly reliable Dali Lama, or just lucky guessing — could be held to yield justified beliefs. And Bonjour himself (1985, p. 107) notes that coherence may be present but not be truth-indicative because of the possibility of multiple coherent sets of beliefs about the same subject, and thus the objection is that coherence can be present but not sufficient for the justification of belief.

Various features across disciplines. We begin by considering a couple of diverse disciplines (or discourses): literary criticism and physics. We assume that these—along with numerous other disciplines—involve epistemic normativity. There are, for example, experts in literary criticism and physics, and there would not be such persons if they were bereft of epistemically normative beliefs. Granting that assumption, the question arises as to whether the epistemically normative beliefs across distinct disciplines are unified simply in virtue of instantiating the same simple and general feature. We argue that they are not.

In the realm of literary criticism, coherence is typically the most significant feature with regard to the epistemic status of beliefs (concerning, say, interpretation). One could well hold that epistemically normative beliefs within literary criticism rest heavily on the practitioner having a coherent body of beliefs concerning a given literary work, genre, or figure. This coherence may be amongst beliefs that are based on individual literary works, correspondences, interviews, background theories, and even interpretations forwarded by other literary critics. What seems clear is that other features — for example, reliability or foundationalist criteria — are often not as dominant or as overriding within the discourse of literary criticism (though they can certainly be part of the normative mix and in some cases even in the normative foreground). Reliability is usually not as dominating because literary criticism is not primarily in the business of tracking phenomena. The discourse does not primarily aim to describe “how the world is” or “how it will be” (though it certainly can involve this; Shakespeare’s alleged “farewell”

intentions behind *The Tempest* is an illustration). The field is instead more directly concerned with the meaning of various texts, genres, etc. In addition, there seems to be lesser room for the type of beliefs espoused by foundationalists, that is, those “basic beliefs” that are justified, but not in virtue of other beliefs. Certainly the literary critic has evidentially basic beliefs for interpretations, for example, beliefs that are basic in the line of reasons the critic gives for an interpretation. Quotations, for example, can be basic to reason-giving in literary criticism and thus are epistemically relevant. But they obtain and confer epistemic normativity mainly in virtue of their coherence with other beliefs that make them relevant, and not in virtue of some non-belief-based justificatory process, for example, appearing on the page to the reader. Hence, we claim also that foundational criteria are (usually) not significantly relevant to the epistemic normativity of beliefs within the realm of literary criticism. Coherence appears to be the overriding feature of epistemic normativity within literary criticism.

The picture of epistemic normativity is different for a field like physics. Even if we allow that the same features of coherence, reliability, and foundational are present, the emphases, we claim, are different. Unlike literary criticism, the epistemic normativity of beliefs in physics is to a significant degree a matter of reliability. Beliefs based on the reliable tracking of phenomena are those beliefs that tend to carry epistemic normativity in the hard sciences. For instance, the having of experimental apparatuses that reliably track proton phenomena within cloud chambers is crucial to epistemic normativity in particle physics, and probably even more crucial than the coherence of our observations of vapor trails in cloud chambers with the theory of protons (and the subsequent theory of how to detect them). It may be reasonably argued, to be sure, that there is a significant role for coherence in scientific belief. We agree that one does not simply “see” a proton in a cloud chamber. The experience is theory-laden — one does not see the direction and charge of the particle by looking at a trace on a photograph. Moreover, the beliefs concerning the experience rely on other beliefs within the theory, and hence the beliefs are not simply foundational either and coherence does have a role. However, we do not want to get embroiled too deeply in debates concerning the epistemology of science as such. What we are minimally claiming with this example is that coherence alone — without reliability — would most often fail to meet the aims of hard science, and that reliability tends to play the leading normative role when it comes to epistemic normativity in the hard sciences.

As for the “soft” sciences, Clifford Geertz (1973, 1988, pp. 1–23) provides us with fairly self-conscious analyses of the issue of how it is that anthropological texts get “epistemic normativity.” He writes, for example, “anthropological texts are themselves interpretations” (1973, p. 15), and this suggests

that features and points mentioned in our discussions of literary criticism and physics may well be in the mix.

The lesson we cull from observing the epistemic practices of diverse disciplines is this: supposing that multiple discourses involve epistemically normative beliefs, then such varied beliefs are often epistemically normative in virtue of different features being epistemically relevant or dominant across the discourses. This is our main point, and at this stage we take it that we have — at the very least — put the ball in the critic's court. And so we move on now to consider two rather predictable criticisms; our treatment of them, especially the second, will lead to a further strengthening of our position.

Sympathetic critics may grant that there is a plurality of features involved in the epistemic normativity of cognition across disciplines but add that the task of epistemology is to formulate second-level principles for when, or to what extent, these features have dominance. Our reply, however, is rather Pyrrhonian and asks in virtue of what feature will this second-level cognition be epistemically normative? Completeness? Convergence (of discourses)? Human flourishing? Non-arbitrariness? Instrumentality?² Here again we would claim that our foregoing argument suggests that there will be no single meta-feature, and note in addition that none of these proposed meta-features needs to “carve things up” in the same way as the others.³

Some critics may respond to the above by arguing for “epistemic austerity,” that is, for the claim that only a very limited scope of our beliefs is epistemically normative. For instance, many, if not most, epistemologists take mathematics and hard science as being paradigms of epistemic normativity, and thus argue that the epistemically distinctive feature in these domains is the basis of all epistemic normativity. Beliefs from the humanities and arts are hence often consigned to the set of beliefs that are merely tastes and opinions, and thus no belief in these realms is epistemically better off than any other — though some may be “more powerful” than others. But we have a response to this challenge that beliefs that fail to meet “the criterion” of mathematics or hard science thereby fail to be epistemically normative.

Various features within a discipline. Our response to epistemic austerity builds on our previous example concerning hard science. We claim that not even *within* hard science is epistemic normativity explicable by a simple and general feature. If scientific belief ranges from common beliefs concerning

²See Putnam (1981, p. 134) for human flourishing, Smart (1963, pp. 1–12) for arbitrariness (in light of Gosse's hypothesis that God created the world as containing fossils and the resulting problem of deciding between hypotheses for which there can be no empirical test), and Moser (1998, p. 14) for instrumentality.

³See Sankey (1997, pp. 138 and 145, n. 2) for a different point involving the problem of criteria regress.

middle-size dry goods to those beliefs concerning the abstract theoretical objects of quantum physics, as Quine (1960), Harré (1986), and others have argued, then we suggest that different features in the normative foreground are apt to make for epistemic normativity at different levels of scientific belief.

If one looks to the beliefs of ordinary persons regarding middle-size dry goods and the associated justification of those beliefs, then one will find that beliefs grounded in non-belief states, for example in sensory perception, are largely epistemically normative. A foundationalist criterion would seem largely to explicate the epistemic normativity of common sense beliefs. However, we claim that this is not predominantly the case for the rest of scientific belief, because more theoretical beliefs are meant to explain and predict — in a deeper manner — sensory phenomena. The epistemology of more theoretical beliefs does not take sensory perception or common beliefs as quite so bedrock. Perceptual illusions, for example, mirages, for example, are explained and predicted by more theoretical beliefs, and explicating the epistemic normativity of these beliefs in turn will involve reference to other features. Perceptual illusions can be explained, to be sure, by other common sense beliefs, but such explanations will be superficial in terms of efficient causality and will show, even at the level of common sense beliefs, a mix between coherence and the foundational feature.

Given that beliefs based on ordinary sensory input are not inviolate within theoretical physical science, it seems to us that many mid-level scientific beliefs — those beliefs captured by accepted scientific theories — carry epistemic normativity largely as a result of their predictive success — a kind of reliable tracking of phenomena — and the coherence involved in their explanation of any given phenomenon. Hence, supposing for simplicity the foundationalist criterion to be largely relevant to and sufficient for the epistemic normativity of ordinary beliefs, the criterion is less relevant to and often insufficient for the epistemic normativity of most mid-level scientific beliefs. Predictive success and coherence, to be sure, can and do play normative roles with respect to ordinary beliefs — it wouldn't be a sane idea for one to invite nonchalantly a neighbor to dinner if one feared gunshots as a response or if one didn't believe that eating is something humans tend to do — but predictive success and coherence are pushed more to the foreground when it comes to the epistemic authority of mid-level scientific beliefs. And while we admit that there can be tensions between predictive success and coherence, we need take no particular stance on how conflicts between the two ought to be resolved. All we argue is that mid-level scientific beliefs generally fail to attain epistemic normativity merely by meeting — in the same manner that low-level beliefs do — the foundationalist criterion.

Finally, at the ethereal level of the most theoretical sciences, for example, areas of high-level physics, predictive success drops out somewhat and is

often inconsequential to epistemic normativity due to the fact that many of the competing beliefs and sets of beliefs (theories) are equally (un)successful in predicting physical phenomena. One need only think about debates over superstring theory and ask how much current predictive success is at stake. At the present time it does not seem as if the acceptance or rejection of superstring theory would make a current difference in the predictive success of physics. Coherence thus tends to reign as the feature most distinctive of those highly speculative scientific beliefs that carry epistemic normativity. What we draw, then, from looking at science is this: if across the range of scientific beliefs epistemic normativity is garnished from the presence of a variety of indicative features, then epistemic austerists are incorrect that paradigms of epistemic normativity are unified by a single feature. In addition, the argument suggests — once again — that the complete systematicity assumption is an unwarranted epistemological assumption.

As a coda to this section, we have learned by our look at literary criticism and hard science that “coherence” and “reliability” each really stand for a cluster of features. In some cases of textual interpretation, for example, what makes for epistemic normativity cases is intratextual coherence whereas in other cases it is coherence between text and background information. As a consequence of this cloaking, proposed *complex* but still relatively simple epistemic features such as “coherent and based on evidence” or “foundherentist” are even further removed from the complete systematicity assumption than many may have thought. To further the point, what counts as evidence is not a uniform matter across disciplines or even within a discipline.

Transition

Some would conclude that our arguments imply that the aim of epistemic unification should be abandoned altogether. Philosophers such as Richard Rorty, Stanley Fish, and Patricia Churchland all advocate or intimate “epistemic eliminativism,” that is, the thesis that traditional epistemology — which includes the project of systematizing justification, warrant, rationality, or what we are calling epistemic normativity — should be abandoned. Concerns like those noted in our examples have, in part, driven epistemic eliminativism.⁴ Indeed, a strand of eliminativist reasoning may be described as follows: if epistemic normativity fails to have a simple unifying feature, then the epistemological project is hopeless, and it should be replaced with some other type of project.⁵ We are not so glum concerning the prospects of

⁴For an example of other concerns driving eliminativism, and especially for the idea that there is a “paradigm shift” in traditional epistemology, see Churchland (1987).

⁵See Haack (1993, pp. 118–202) for thorough discussions and refutations of the most pernicious forms of eliminativism.

unifying epistemology; we are just not as hopeful — or perhaps just more willing to throw out supposed constraints on what makes for an epistemological theory — as those epistemologists who continue to adhere to the ideally-complete systematicity assumption.

If we look further into the above discussion concerning the structure of normativity in ethics, we see that there is a type of unification of moral normativity that does not hold that moral normativity must be grounded in a simple unifying feature. And yet this view does not abandon altogether the aim of systematizing morality. Here we are thinking of moral pluralism, especially as it is articulated by W.D. Ross (1930), who defends the view that moral normativity cannot be reduced to or captured by a simple and universally relevant feature. Ross argues that neither monistic consequentialism nor monistic deontology are free from counter-examples, and concludes that “it is more important that our theory fit the facts than it be simple” (p. 15). As a result, Ross abandons monism for pluralism in ethics because he believes that examples both for and against individual features prove that certain features are not always morally overriding. But Ross further believes that the examples indicate the “normal” moral relevance of the features. He thus proposes that the features are to be captured by rules, but rules that contain ineliminable *ceteris paribus* clauses due to the defeasibility of the features on which the moral properties of actions normally supervene. Feature pluralism and non-exceptionless rules thus make up Ross’s system of *prima facie* duties. For similar reasons, that is, examples that indicate both the overridability and normality of certain epistemic features, we think that a similar “pliant” structure, and a corresponding “partial” unification, are appropriate to the project of systematizing epistemic normativity.

Pliant Epistemic Normativity

The foregoing arguments against the two assumptions are not merely of negative import: they suggest an alternative view on the structure of epistemic normativity. The arguments undermine the idea that the full range of epistemic normativity can be completely expressed in terms of any theory that mirrors a rigid structure, but the arguments leave open the possibility of a sort of partial unification via a theory that mirrors a pliant structure. A pliant structure of epistemic normativity is specifiable as a fairly simple, fairly complete system, and it is a suitable underpinning, for example, for epistemological theorizing that aims after “defeasibly-partial epistemic unification,” that is, the somewhat incomplete accounting of the full range of epistemic normativity in terms of defeasible principles and various defeasible, general, and relatively simple features. The examples from particularism, to be sure, seem to militate against even this possibility, but we see those examples and

radical particularism itself as being tempered by the other arguments. Particularity may well be ineradicable from ethical and epistemic contexts, but that does not preclude them from also exhibiting patterns of some sort, for example, patterns that allow exceptions. And so we can admit of a middle position that takes the structure of epistemic normativity to be neither rigid nor amorphous. We call this position “pliant epistemic normativity,” and the heart of it is that the structure of epistemic normativity is specifiable as a fairly simple, fairly complete system.

The approach to epistemological matters in terms of “rigid,” “amorphous,” and “pliant” structures is admittedly not the most familiar of tacks, and so a quick analogy might help to illustrate what we are suggesting for epistemological theorizing with the idea of pliant epistemic normativity. One of the components in Kant’s moral philosophy is a claim concerning, in effect, the “structure” of morality. The claim is that the correct morality will be comprised of a set of categorical imperatives—whatever their content turns out to be. “One ought to keep promises” is an example of a categorical imperative, and so is “One ought not to murder.” Kant holds that sentences of this form, that is, universally-binding prescriptions, are of the kind that is appropriate to moral philosophy. One thus gets the image of an overall structure onto which the content of morality is to be mapped. In an analogous way, we are developing the claim that the correct epistemological theory — given the project of unifying epistemic normativity — will be comprised of a set of defeasible features captured in principles with ineliminable *ceteris paribus* clauses — whatever their content turns out to be.

We should note here that we are urging that pliant epistemic normativity is even more pliant than what is (implicitly) suggested by certain epistemological theories. For instance, Goldman’s (1988, pp. 51–69) distinction between weak and strong justification has the appearance of mirroring pliant epistemic normativity. Strong justification, according to Goldman, is externalist in nature and is simply a matter of whether or not a belief is the product of a reliable cognitive process. Weak justification, on the other hand, is internalist and is based merely on whether epistemic agents have reasons that they believe are sufficient to endow epistemic normativity. Now, even though Goldman grants that weak justification is a form of epistemic normativity and that this adds a measure of pliability to his account, we think it is not pliant enough. We argue, against Goldman, that the features in virtue of which beliefs are epistemically normative — regardless of the internalist/externalist debate — are multiple. There is no single feature, in other words, that is always epistemically relevant or overriding — even if that feature is “being the product of a reliable process.” We maintain that the history of epistemological debate and our above arguments indicate that the grounds of strong justification are in fact plural and defeasible. This is not true of

Goldman's full formulation of epistemic normativity. For Goldman, all internalist features are defeasible, but the externalist feature of "being the product of a reliable belief-forming mechanism" is not defeasible and it is always epistemically overriding. This conjunction drives his claim that for a belief to be "fully justified" is for it to be the product of a reliable belief-forming process *and* to be defensible with reasons that are appropriately related to that process. In opposition, we are claiming that full justification — epistemic normativity, in our words — is itself pliant, and that the pliability is not solely a matter of differences between internal and external justification.

To return to our view, the defeasibly-pliant structure suggests that the correct epistemological theory must not only capture those features that are typically relevant but also do so in a way that is flexible enough to account for those cases where any given feature, or combination of features, is not relevant or not overriding. The best way to codify and yet reflect such features is with epistemic principles that contain ineliminable *ceteris paribus* clauses. For example, one might argue that beliefs grounded in appropriate non-belief states are epistemically normative, but the recognition of pliable structure would lead one to formulate the epistemic principle as follows: "If a belief is based on unobstructed sensory perception, then *ceteris paribus* the belief is epistemically normative." The advantage of this formulation is that it not only incorporates the fact that beliefs of this type are not always epistemically normative — because *ceteris* is not always *paribus* (as with certain illusions, in which cases coherence may override the foundationalist feature) — but it also captures what *quite often is relevant to* the epistemic normativity of certain beliefs. To deny that epistemic principles are exceptionless is not to disallow that they are generalities of some other sort. As for what the features are, we are here leaving that question open, but we do believe that the features traditionally defended in epistemology are most likely to be included in the set of features that are indeed distinctive of epistemic normativity.

We close by offering a response to an anticipated question, and in so doing underline an important implication of the defeasibly-pliant structure for epistemological theorizing. We would claim, in response to the question of how to adjudicate between competing features and principles, that however this is to be done, no single feature or compact set of exceptionless principles will do the trick. That is our point. The paradigm of complete unification would have us ascend to some meta-level feature and principle, but the paradigm of partial unification would have us focus on other strategies altogether.

Conclusion

Ideally-complete epistemic unification, as we understand it, requires that the structure of epistemic normativity be specifiable as a maximally simple,

maximally compact system. The ideal, across-the-board unification of epistemically normative cognition, in other words, requires that certain structural assumptions are true. But the complete systematicity and complete codifiability assumptions — the required structural assumptions — are not true, and thus ideally-complete epistemological unification is not possible. Our case can be summarized as follows: (1) Frame-problem cognition points to a codification of epistemic normativity that stops short of ideally-complete codification by a compact set of exceptionless epistemic principles. (2) Frame-problem cognition points to a partial codification by generalities of a non-exceptionless kind. (3) Beliefs from across disciplines are often epistemically normative in virtue of different features being in the normative foreground, for example, the epistemic normativity of beliefs in literary criticism as opposed to physics. (4) Beliefs from within the continuum of a single discipline are often epistemically normative in virtue of different features in the normative foreground, for example, the epistemic normativity of beliefs at different levels of scientific discourse. (5) Points 1–4 do not entail epistemological eliminativism, however, for if epistemic normativity has a pliant structure, then a qualified sense of epistemological unification is possible, and thus there is a point to this kind of epistemological theorizing. Thus: (6) Epistemic patterns of a defeasible sort may well be found across contexts.

We have not, we should like to re-emphasize, touched on entirely different conceptions of epistemological theorizing, and thus even if a pliant structure were not possible, the upshot would still not necessarily be epistemic eliminativism. But to continue the review, the pliant structure of epistemic normativity points to a positive aim with regard to epistemological theorizing. *The pliant structure suggests that epistemic theorizing should aim to make explicit the features and principles that partially unify epistemic normativity.* The achievement of this aim would result in the defeasibly-partial unification of epistemology, and this may well be the best that can be done. But defeasibly-partial unification is still better than eliminativism, for the former allows for the generality that helps to make for systematic understanding, and thus it could be endorsed as a legitimate aim for epistemological theorizing. One can thus be hopeful despite despair over full-blown unification. One can be impressed by the above arguments adverting to the contextuality that epistemology must confront, but one need not be defeated by them.⁶

We are not alone in offering this proposal concerning the aim of epistemological theorizing. A one-time seeker of complete epistemological unifica-

⁶We are thus not overly impressed by Rorty's (1989, p. 73) "irony," viz., that we make judgments even though we are not justified. Why assume we are justified in assuming relativism? At least one of us, rather, openly laments the fact that we in the contemporary world do not make judgments even though we are justified.

tion, William Alston, is now writing in a different key. The overall tone of his "Epistemic Desiderata" pleases our ear:

... we should abandon the idea that there is a unique something or other properly called "epistemic justification." Having done so we will be free to recognize and investigate a number of different ways in which beliefs can be better or worse from an epistemic point of view. (1993, p. 527)

We endorse this conclusion, but we have here presented additional and more fundamental reasons than Alston's for it. Alston's reason is that *the persistence of disputes among epistemologists* concerning epistemic justification leads to the suspicion that there is not enough commonality in pre-theoretical understandings of epistemic justification to suppose that there is some uniquely identifiable object called epistemic justification (pp. 532, 534). We have gone "underneath" the fact of persistent dispute and explained why the lack of pre-theoretical commonality is inevitable. Moreover, we are more hesitant than Alston seems to be with respect to "dropping the question of justification of belief altogether" (p. 527); by allowing for the conditions under which beliefs are justified to be multiple and defeasible we are going between the horns of the dilemma.⁷

References

- Alston, W. (1993). Epistemic desiderata. *Philosophy and Phenomenological Research*, LIII(3), 527–551.
- Audi, R. (1993). *The structure of justification*. Cambridge: Cambridge University Press.
- BonJour, L. (1985). *The structure of empirical knowledge*. Cambridge: Harvard University Press.
- BonJour, L. (1997). Haack on experience and justification. *Synthese*, 112, 13–23.
- Churchland, P. (1987). Epistemology in the age of neuroscience. *Journal of Philosophy*, xxxiv, 544–553.
- Geertz, C. (1973). *The interpretation of cultures*. New York: Basic Books.
- Geertz, C. (1988). *Works and lives: Anthropologist as author*. Stanford: Stanford University Press.
- Goldman, A. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge* (pp. 1–23). Dordrecht: Riedel.
- Goldman, A. (1988). Strong and weak justification. In J. Tomberlin (Ed.), *Philosophical perspectives, volume 2* (pp. 51–69). Atascadero, California: Ridgeview.
- Haack, S. (1993). *Evidence and inquiry: Towards a reconstruction in epistemology*. New York: Basil Blackwell.
- Harré, R. (1986). *Varieties of realism*. New York: Basil Blackwell.

⁷Our overall conclusion does not imply that we are particularly anxious to be compared with virtue epistemologists. If epistemic normativity is indeed of a pliant structure, this does not necessitate that epistemological theorizing be wed to reliabilism or reject all degrees of codifiability. It is rather still possible that epistemological theorizing should seek to formulate in general terms the features and principles that are distinctive of epistemic normativity. Moreover, our overall conclusion fits nicely with epistemological contextualism.

- Haselager, W. (1997). *Cognitive science and folk psychology*. Thousand Oaks, California: Sage.
- Henderson, D., and Horgan, T. (2000). Iceberg epistemology. *Philosophy and Phenomenological Research*, LXI, 497–535.
- Horgan, T., and Tienson, J. (1996). *Connectionism and the philosophy of psychology*. Cambridge, Massachusetts: The MIT Press.
- Moser, P. (1998, November). *Epistemology today*. Paper presented at Northern Illinois University Graduate School Colloquium, DeKalb, Illinois.
- Putnam, H. (1981). *Reason, truth and history*. Cambridge: Cambridge University Press.
- Putnam, H. (1983). *Realism and reason: Philosophical papers, volume 3*. New York: Cambridge University Press.
- Pylyshyn, Z. (Ed.). (1987). *The robot's dilemma: The frame problem in artificial intelligence*. New Jersey: Ablex.
- Quine, W. (1960). *Word and object*. Cambridge, Massachusetts: The MIT Press.
- Rorty, R. (1989). *Contingency, irony, and solidarity*. Cambridge: Cambridge University Press.
- Ross, W.D. (1930). *The right and the good*. New York: Oxford University Press.
- Sankey, H. (1997). *Rationality, relativism, and incommensurability*. Aldershot, England: Ashgate.
- Smart, J.J.C. (1963). *Philosophy and scientific realism*. London: Routledge and Kegan, Paul.