

**WHAT KNOWLEDGE IS AND WHAT IT OUGHT TO BE:
FEMINIST VALUES AND NORMATIVE EPISTEMOLOGY**

Sally Haslanger
Department of Linguistics and Philosophy, MIT

I. Introduction¹

Much of contemporary analytic epistemology has been concerned with the semantics of claims to know: What are the truth conditions of claims of the form *S* knows that *p*? With some notable exceptions, feminist epistemologists have not taken up this project—at least not in this form—so for those who are engaged in mainstream epistemology it may seem tempting to think that what feminists are doing is not relevant to their concerns, and to ignore feminist work as addressing a different set of issues.² Although I think it is right that a lot of feminist epistemology is addressing different issues, this response does not take into account that a significant amount of feminist writing explicitly undertakes to critique the mainstream epistemological questions; it is not simply that feminists are interested in something else, but that they have principled reasons for *not* engaging the issues as standardly framed.

My interest in this paper, however, is not in evaluating the feminist challenges to the search for the truth conditions for knowledge claims. Although I am sympathetic with the complaint that there are many other epistemological topics that mainstream epistemology could and should consider, I am not convinced that this project itself is misconceived or irretrievably sexist or androcentric. I do think, however, that there are problems with the ways that philosophers have undertaken to provide an analysis of knowledge, and the problems suggest that an alternative approach informed by feminist concerns is desirable. My goal in this paper is to suggest a way of approaching the task of specifying the truth conditions for knowledge, that (hopefully) will make clear how a broad range of feminist work that is often deemed irrelevant to the *philosophical* inquiry into knowledge is in fact highly relevant.

II. The Questions

Questions of the form, "What is X?" (or What is it to be an X?) are often used to demand an articulation or clarification of the concept of X: What is the (ordinary?) concept of *knowledge*?³ What is the proper analysis of the concept? I'll call this kind of project a *conceptual* investigation into X. Traditionally conceptual projects were treated as wholly apriori affairs, but contemporary efforts sometimes allow a degree of sociological or anthropological investigation in considering the variety of uses of the term or concept in question.

In contrast to the conceptual project, a *descriptive* or *naturalistic* project concerning X is not primarily concerned with exploring the nuances of our concepts (or anyone else's for that matter); it focuses instead on their (purported) extension, i.e., the things that (purportedly) fall under the concept. Here, the task is to develop potentially more accurate concepts through careful consideration of the phenomena; this is achieved by establishing empirical or quasi-empirical generalizations about the domain in question. Paradigm descriptive projects occur in the natural sciences where the goal of understanding, e.g., what water is, is not to analyze our ordinary concept of water, but to offer an account based in an empirical study of the relevant phenomena. In the case at hand projects in naturalized epistemology seek to answer questions such as "What is knowledge?" through an aposteriori investigation of what we normally take to be paradigm instances.

In recent years there have been two different kinds of naturalistic projects in epistemology. The first and more radical form assimilates epistemology to psychology (or sometimes sociology): the idea is to take our ordinary knowledge attributions to fix the reference of our epistemic terms, and then to undertake an (aposteriori) investigation of the natural (or social) kinds that are (allegedly) being referred to. The question is what, if anything, do those things that normally get called knowledge have in common? Do they deserve to be considered a kind—are they a unified collection? And if so, what is the basis for their unity? An alternative and recently more popular version makes explicit room for a normative component in knowledge by seeing the project as an investigation into the supervenience base for our ordinary epistemic evaluations.⁴ Assuming that the normative supervenes on the non-normative, and that epistemology is normative, the question is: On what non-normative (physical, psychological, or social) facts does knowledge supervene? Both kinds of naturalizing approach begin the inquiry with pre-theoretic intuitions about cases that "fix the reference" of the term, and yet the resulting accounts often demonstrate the need for conceptual revision and can even serve to debunk the ordinary concept entirely (e.g., if the extension of our term 'knowledge' is not a natural kind, or if nothing is found to provide the supervenience base).

In practice conceptual and descriptive projects can't be kept entirely separate, for each typically borrows substantially from each other. Conceptual projects depend upon a careful consideration of "normal" or paradigm cases and descriptive projects can provide the detailed accounts of them needed; in turn, descrip-

tive projects require a rough specification of the boundaries of the phenomenon to be investigated and depend on conceptual projects to circumscribe what sorts of cases are at issue. In fact, the difference in these kinds of projects might be best taken to be a matter of emphasis. But the two sorts of project differ importantly in both their guiding questions and conditions of success: a conceptual project is concerned to specify, of all the candidate concepts, which concept of knowledge is *ours*; a descriptive project is concerned to specify, of all the candidate (natural) kinds, which we are referring to by the use of our epistemic terms. Although conceptual and descriptive projects are the most common (contemporary) approaches to questions of the form "What is X?" they are not the only ones; in fact, we'll consider a third approach later in this essay.

III. Knowledge and "Everyday Practices"

There is no doubt that much of modern and contemporary epistemology has been framed as an effort to respond to epistemological skepticism. The main anti-skeptical strategy of 20th century epistemology has been to challenge the skeptic's conception of knowledge by taking our actual knowledge practices as providing definitive cases of knowledge. Beginning with the assumption that we do have knowledge in at least some ordinary contexts, the task is to use these contexts as a basis for articulating a conception of knowledge—possibly a naturalized conception—that might properly be called "ours," and that also rules out the skeptical hypothesis. The broad suggestion is that we should reject a "transcendental" epistemology that imposes conditions on knowledge that presume a standpoint outside of our practices, and should instead pursue what we might call an "immanent" epistemology that undertakes to elucidate the conditions on knowledge embedded in our everyday language, thought, and action.⁵ Although the skeptic purports to be using the term 'know' as we do in claiming that we do not know there is an external world, attention to our use of epistemic terms shows that the skeptic is in fact employing a different concept whose conditions for use are not ours. Hence the skeptical challenge does not undermine our ordinary claims to know.

Of course even if there is broad agreement on the strategy of "immanent epistemology", controversy remains, for it is unclear what conception of knowledge is embedded in our "ordinary practices". Epistemologists, whether engaged in a conceptual or naturalistic project, have undertaken to uncover the "embedded" concept. For example, ordinary language philosophers have attempted to elucidate a nuanced analysis of our concept through a more or less apriori investigation. And radical naturalizers use our everyday attributions as the starting point for their investigation into the natural kinds that we are (allegedly) referring to. With this focus on our ordinary knowledge attributions, defenders of skepticism have, in turn, argued that in fact the conditions for knowledge embedded in our everyday practices are the skeptic's after all. But even if this is so, at this stage of the debate any theorizer—skeptical or not—who wants to maintain

that the account they offer is an account of "our" concept of knowledge is committed to showing how it is to be found in our practices.

IV. Feminism and the Pragmatics of Knowledge

On the face of it, then, it would seem that feminist work examining and critiquing our everyday knowledge practices would be extremely valuable to anyone undertaking an analysis of knowledge. If the goal is to offer an account of "our" concept of knowledge, then it is an important question whether the concept "embedded" in our practices is sexist, androcentric, or otherwise politically problematic.

In fact, feminists have documented in impressive detail that our actual practices of knowledge attribution are both sexist and androcentric. Consider three kinds of questions about our ordinary practices (I raise them here for gender, but they can also be raised for race and class):

- i) Is an individual's gender relevant to whether he or she is likely to claim and/or to be attributed knowledge? And is one's gender relevant to the *domain* in which one is likely to claim and/or be attributed knowledge?
- ii) Are the methods that are likely to be counted as knowledge-producing more often associated with men than women? Can the hierarchy of kinds of belief/method be justified on epistemic grounds or does the hierarchy reflect gender bias?
- iii) Are the conventions of authorizing certain individuals as knowers and the social rituals that accompany such authorization, e.g., rituals involved in deferring to those authorized and in challenging authority, problematically sexist (or problematic in other ways)? Do these conventions have problematic effects on the workings of knowledge communities, e.g., do they exclude women and protect ideological views from being challenged? Do they foster attitudes towards the natural world and towards other people that are androcentric and morally questionable?

In addressing these questions, feminists have accumulated substantial evidence that our actual knowledge attributions and practices of authorization privilege men and help sustain sexist and racist institutions. It is not essential to my project here to make the case that there is sexism in our everyday epistemic practices, for my concerns are more methodological. I would hope that it is obvious to anyone who has reflected for even a moment on their own behavior and the behavior of those around them that cognitive authority is not taken or granted in gender-neutral ways, and the *prima facie* plausibility of that claim is enough to raise the questions I want to address. But it may also be helpful to indicate briefly some of the main areas in which feminists have documented concerns⁶ (again some of these points are directly parallel to ones that can be made concerning race and class):

Sexism and sex-stereotyping in attributions of knowledge:

Refereeing: acceptance rate of papers by women increases when procedures are implemented to prevent the referee from knowing the sex of the author.

Classroom climate: women/girls and men/boys are asked different sorts of questions with different sorts of follow-up; men/boys are more often assumed to be capable in learning the subject and women/girls are not.

"Masculine" and "feminine" coding of methods and fields:

Quantitative "hard" research is coded as masculine and is considered more important and more valuable; but these privileged methods are not uniformly successful and often what is considered "feminine" research is more effective and/or addresses different though equally valuable domains of inquiry.

Entrenchment of sexist ideology:

Theories that affirm the naturalness of current sex roles (and other ideological expectations) are more quickly endorsed.

Use of sexist/gender metaphors in understanding non-gendered phenomena reinforce the idea that sexist social arrangements are "natural".

Outright sexism in research communities.

Feminist discussions of these phenomena are tremendously rich and suggestive. However, the standard reply to taking these feminist studies to be philosophically illuminating is that it is "just sociology, not philosophy"—feminism tells us a lot about the sexism in our communities, but not much about our concept of knowledge. Stated as simply as this, however, the reply doesn't have much force coming from an epistemologist who favors an "immanent" strategy, for if we are trying to discern the concept embedded in our practices, "sociological" information about those practices should be relevant.

However, even though the simple reply can be dismissed, more needs to be said to link the sexism in our practices with the truth conditions for knowledge claims: How exactly should one go about reading our "embedded" concepts off our practices? Let me use a somewhat exaggerated example to demonstrate the problem. Suppose that there is substantial and systematic sexism in our attributions of knowledge. Should we conclude from this that "our" concept of knowledge is one that requires the knower to be male (or masculine)? For example, should we analyze the concept embedded in "our" practice along these lines:

S knows that p iff S is justified in believing p, p is true, and S is male (or S is in some relevant respects masculine).

Or should we provide an account with different justification conditions for men and for women, e.g., requiring women to have greater justification than men in order to count as having knowledge? These suggestions are implausible (and I don't think it has been part of feminist epistemology to defend anything like this). The problem is that there are several ways to account for sexist attributions of knowledge other than claiming that the concept of knowledge being employed has the sexism built in to its truth conditions. E.g., one might claim that the conditions for knowledge are gender neutral, but that background sexist beliefs lead people to believe that men are more likely to satisfy the conditions than women.⁷ One way to develop this explanation would be to draw on a distinction between linguistic or conceptual competence and performance.⁸ It is not unusual for ordinary and perfectly competent speakers of the language to get things repeatedly wrong due to systematic distortions present in the context; perhaps we should understand pervasive sexism as one of those systematic distortions that prevents us from making correct epistemic evaluations even by our own lights.

Alternatively, one could resist the charge that the sexism in our practices reveals a gendered concept of knowledge by arguing that in making knowledge attributions we are doing more than asserting that someone meets the conditions for knowledge, i.e., the utterance conditions for knowledge attributions should be distinguished from the truth conditions.⁹ If so, then our practices involving epistemic utterances may have us differentiate men and women not because the concept of knowledge employed in these utterances is somehow gendered, but because what we are doing with these utterances (besides asserting the knowledge attribution) is politically problematic.

To take this latter route is to allow that claims to know do more than assert propositions. Knowledge claims may well have propositional content, but expressing that propositional content is not the sole function of our speech act in claiming to know. As Austin so vividly puts it:

...saying, 'I know'...is *not* saying 'I have performed an especially striking feat of cognition, superior, in the same scale as believing and being sure, even to being merely quite sure': for there *is* nothing in that scale superior to being quite sure. Just as promising is not something superior, in the same scale as hoping and intending, even to merely fully intending; for there *is* nothing in that scale superior to fully intending. When I say, 'I know', I *give* others my word: I *give* others *my authority for saying* that 'S is P'.¹⁰

He later continues:

If you say you know something, the most immediate challenge takes the form of asking, 'Are you in a position to know?': that is, you must undertake to show not merely that you are sure of it, but that it is within your cognizance.¹¹

Austin's reflections suggest that in making first person claims to know one is not (or not simply) reporting a psychological or cognitive state; one is performing a

certain kind of socially meaningful act: among other things, one is claiming epistemic authority, an authority that may be "given" to others, and that may, in turn, be challenged.¹² Admittedly, one might interpret Austin as reading the truth conditions for knowledge claims directly off of our practices of knowledge attribution, but his own distinction between the locutionary and illocutionary force of an utterance provides room for a wedge between our epistemic practices and our epistemic concepts.

As I understand it, a lot of feminist work in epistemology focuses on the illocutionary force of knowledge claims—what's done, how it is done, the rituals and conventions that govern the distribution of epistemic power—rather than specifically addressing the question of truth conditions. As suggested above, the question isn't only *who* is authorized, and *what* methods are authorized, but how the rituals of authorization create and sustain self-affirming ideological communities. But acknowledging that our epistemic practices are mechanisms for the distribution of power and authority still leaves open a difficult question: what is the relationship between the (plausibly problematic) conditions of utterance for knowledge attributions and the truth conditions? Couldn't one reasonably maintain that feminist theorists should engage in a critique of our epistemic practices, and even allow further that this is an important part of epistemology, but still claim that in spite of a devastating critique of our practices our basic concept of knowledge remains intact?

V. Reflective Equilibrium and the Search for "Our" Concepts.

Note that both of the replies just sketched—one drawing on the distinction between performance and competence, the other between utterance conditions and truth conditions—assume that if we are careful about how we proceed, then there is some way to home in on a concept of knowledge that can rightly be considered "ours" through an examination of our ordinary epistemic practices. But what we need is some way to sort our attributions of knowledge into those that are properly indicative of our concept, and those that are not. What is the best way to do this? The standard procedure is to employ some form of reflective equilibrium: (roughly) consider the range of typical applications of the term in question and the generally agreed upon ("pre-theoretic") principles thought to govern its use, and determine the set of conditions that best accommodate both the applications and principles (allowing that cases and principles can be weighted according to centrality or importance).

This use of the method of reflective equilibrium, it might be thought, does well to set a standard for competence that allows performance errors: once examined closely the mistaken performances will be shown not to accord with the principles we endorse. And it might be possible to develop a sophisticated version of a reflective equilibrium test in order to distinguish truth conditions from utterance conditions. But there are compelling reasons, I think, not to rely on the method of reflective equilibrium (or at least "narrow" reflective equilibrium) as

a basis for defining "our" concept of knowledge, reasons that should lead us to consider another way of pursuing an "immanent" approach to epistemology.

My doubts overlap substantially with those articulated by Stephen Stich in his critique of (what he calls) "analytic epistemology".¹³ Stich's critique does not arise from feminist discussion of our epistemic practices, but in his case as in ours, the question is what to make philosophically of the messy reality of our epistemic and doxastic lives. As Stich points out, employing the method of reflective equilibrium in order to find the target concept in our practice presupposes (i) that there is no more than one such notion embedded in our practices, (ii) that there are general principles that govern our use of the notion or notions, (iii) that there are effective ways of distinguishing principles that constitute meaning and ones that are part of our background "folk" theories (i.e., analytic from synthetic principles), and (iv) that our practices and the concept in question are each coherent.¹⁴ But most of these assumptions are highly questionable, especially for our epistemic notions.

But the second, and more important, question is why we should place so much weight on "our" concept in the first place. If we allow (which, given both the depth of our capacity for cognitive error, and the depth of our sexism, seems reasonable) that our practices might be systematically misguided; and if we take the primary task of epistemology to be a normative investigation into knowledge—one investigating how we *ought to* reason, on what basis we *ought to* form beliefs, and more generally what is epistemically *valuable*—then there is something peculiar about pursuing an "immanent" strategy in epistemology that undertakes simply to describe "our" concept, or to discover the (natural?) kind we ordinarily refer to. Normative epistemology certainly has much to learn from close attention to the ways we proceed epistemically, but to suppose that what we value epistemically is what we *ought to* value epistemically is to leave the normative part of normative epistemology undone.¹⁵

More specifically, the reflective equilibrium strategy of using our intuitions about principles to check our intuitions about cases and vice versa is not a plausible way to sift out sexist assumptions if we have reason to think that our everyday intuitions about cases and principles are interdependent, and both subject to pernicious background influences. For example, consider the question: can emotions count as (defeasible) evidence for a claim? Whatever pressures there are to think that they can't would seem to apply both to the specific judgments we make about cases (of course Susan is not justified in believing *p*, she is just in a fit of rage...), and the principles we affirm (to be justified in believing *p* one must have engaged in critical reflection upon one's beliefs...and so those who believe *p* in a moment of rage cannot be justified).

Note that I'm not suggesting that we should retreat from an "immanent" approach to epistemology altogether, if we understand an immanent approach to be one that undertakes to provide an analysis of knowledge *informed by* our "everyday language, thought, and action". In fact, it is through reflection on our everyday practices of knowledge attribution that I think we can confirm

that our epistemic vocabulary functions normatively.¹⁶ Nor am I suggesting that we abjure all use of the method of reflective equilibrium, for it may be that our best method for determining what is valuable is to engage in a wide reflective equilibrium—where what's involved is a broad critical reflection on one's beliefs, attitudes, and practices, to determine what combination, if any, is reflectively endorsable.¹⁷ My concern is that because we employ our epistemic vocabulary to evaluate each other cognitively, we must undertake a normative inquiry into what is epistemically valuable, and not simply assume that "our" ordinary concept of knowledge—even when modified by recognized experts—captures what we should value. The approach I favor, then might reasonably be considered a form of immanent epistemology, but a *critical* or *normative* immanent epistemology.

VI. Babies and Bath Water

Let's step back for a moment to consider where we stand. So far I've suggested that we—meaning to include both feminist and non-feminist epistemologists—may have good reasons for resisting a non-critical/non-normative approach to epistemology. A purely "descriptive" approach to the analysis of knowledge, I've argued, either ignores the normative question of what epistemic concepts we ought to employ, or assumes implausibly that the epistemic concepts we do employ are the ones we ought to.

But from a feminist point of view it may seem that I've thrown the baby out with the bath water: Wasn't my main point earlier that immanent epistemology cannot afford to ignore feminist research on the role of gender in our epistemic practices because it is committed to articulating the concept "embedded" in those practices? If our goal is no longer to explicate our actual concept of knowledge, how does the feminist research remain relevant? How does the discussion so far help us see how feminist research matters to epistemology at all?

My answer, I'm afraid, is going to be very programmatic. As I see it, the best way of going about a project of normative epistemology is first to consider what the point is in having a concept of knowledge: what work does it, or (better) could it, do for us? and second, to consider what concept would best accomplish this work. To frame the project this way is to employ a different approach to answering the question, "What is knowledge?" (or more generally, "What is X?") than either the conceptual or the descriptive approaches outlined above. I'll refer to this third sort of project as an *analytical* approach.¹⁸

On an analytical approach the task is not simply to explicate our ordinary concept of X; nor is it to discover what those things we normally take to fall under the concept have in common; instead we ask what our purpose is in having the concept of X, whether this purpose is well-conceived, and what concept (or concepts) would serve our well-conceived purpose(s)—assuming there to be at least one—best.¹⁹ Like the descriptive approach, this approach is quite comfortable

with the result that we must revise—perhaps even radically—our ordinary concepts and classifications of things.

Some analytical projects are oriented towards theoretical concepts: the concept *X* is explicitly introduced or adopted as a theoretical tool within a larger inquiry, where the emphasis in determining the content of the concept is placed on the theoretical role it is being asked to play.²⁰ But an analytical approach is also possible in exploring non-(or less-)theoretical concepts if we are willing accept an answer to the question “What is *X*?” that does not exactly capture our intuitive concept of *X*, but instead offers a neighboring concept that serves our legitimate and well-conceived purposes better than the ordinary one.

So on an analytical approach, the specifically epistemic questions “What is knowledge?” or “What is objectivity?” require us to consider what work we want these concepts to do for us; why do we need them at all? The responsibility is both to investigate our purposes in having them, and then to define them in a way that best meets our legitimate purposes. In doing so we will want to be responsive to ordinary usage (and to aspects of both the connotation and extension of the terms). However, there is also a stipulative element to the project: *this* is the phenomenon we need to think about; let us use the term ‘knowledge’ to refer to it.²¹ In short, on this approach, it is up to us to decide what to count as knowledge, and more radically, whether there is anything in our current usage of the term that compels us to carry on with any distinctions along those lines at all.²²

It is plausible to think that there are several different purposes being served by our epistemic practices, and we need to ask what those purposes are, how they are related, and to what extent they are legitimate.²³ Feminist work is relevant to normative epistemology because such work contributes to the exploration of what our purposes are, what they could be, and what they ought to be, in employing an epistemic framework. Feminist work also provides creative alternatives to existing conceptual frameworks for serving our legitimate purposes; given the revisionary potential of normative epistemology, the broadening of our conceptual resources offered by feminists should be welcome.

These suggestions are too general to be convincing, however, so it is probably better to describe the task less blandly and more politically. Some of our purposes in having an epistemic framework are likely to be very basic animal purposes—we need to have some relatively reliable information to help us get around in the world, we need to be able to adjudicate when other animals have information we can effectively use, etc. Cognitively, we are both limited and empowered by our animal embodiment. But knowledge is not, and has never been thought to be, simply true belief, and human knowledge communities and the norms that define them do more than facilitate the gathering and exchange of information: they draw lines of authority and power, they mediate each person’s relationship with herself (in defining conditions for self-knowledge), they circumscribe common ground for public debate and the basis for public policy (and much more). To decide what is epistemically valuable we need to decide what kind of knowledge community is desirable, and this can’t help but involve polit-

ical priorities and political choices. Feminists have much to contribute in considering such priorities and choices.

VII. Constitutive Epistemic Values

One might object to the picture I've started to sketch by insisting that there are narrowly defined epistemic goals that should dictate our epistemic commitments, and consideration of political goals is not appropriate in determining what constitutes knowledge. To develop this point we can borrow a distinction drawn by Helen Longino between the *constitutive* and *contextual* values of our epistemic practices: the constitutive values of a practice are those that constitute the goal or end of that practice, or are necessary means towards those ends; in contrast, contextual values are those present in the context where knowledge is sought, discovered, attributed, denied, forgotten, etc., but that do not define what the practice is, or what makes it an *epistemic* practice.²⁴ The thought then is that communities can decide what weight to give epistemic values such as truth, objectivity, coherence, etc., compared to other values (e.g., some communities may decide that objectivity is not as important as solidarity)²⁵, and this is certainly a political matter, but the *epistemic* status of such values as truth and coherence is not up for political negotiation, for they are the constitutive values of anything that could be considered an epistemic practice in the first place. On this view, feminist debate about the politics of knowledge is not relevant to determining what knowledge *is*, or better *ought to be*, for what epistemic concepts we ought to employ should be determined by what is epistemically valuable, and not by political concerns.

There is something right about this complaint, but as it stands it is inadequate, for defining what's *epistemically* valuable in terms of the constitutive values of our epistemic practices just pushes the normative question back. Given the critique of non-normative/non-critical epistemology we've just been through, the question ought immediately to arise: what recommends our epistemic practices as opposed to some others? Why care about the epistemic values embedded in our actual practices, especially if we have reason to be critical of those practices?

It might be argued, however, that this reply misunderstands the objection: the constitutive values/goals of knowledge are not to be understood by considering the goals of our actual practices, but by reflection on the attitude of knowing (and perhaps the idealized practices that are required of those who aim for knowledge). What we're looking for are the constitutive values of knowledge itself, not our current knowledge practices. To get a handle on this it is useful to begin with belief: our epistemic evaluations involving the concept of knowledge look to be concerned with the question of what it takes to be exemplary in believing something. And beliefs are just the sorts of things that bring with them their standards of evaluation: a belief is correct if true; moreover, if a state isn't in the business of being true, then it isn't belief.

So consider the claim that truth is the constitutive goal of belief, that it is essential to anything that might count as a belief that it "aims at" the truth.²⁶ This

a claim about the nature of belief: for a psychological state to qualify as a belief it must represent its content as true, and in addition, the belief is correct only if the content is true.²⁷ (Contrast, for example, believing *p* from imagining *p*, or considering *p*. Many cognitive states represent their contents as true, but belief is distinctive in that the point of belief is to represent *p* as true only if it is.²⁸ Likewise a belief is correct or apt only if it is true; this is not the case of imaginings or considerings.) So if we are looking for what we ought to value in the cognitive domain, and if we frame these questions in terms of what *believers* ought to value, how *believers* ought to proceed, then it would be paradoxical to deny the value of truth. To be a successful believer is to represent the world accurately. (This is not because we must assume that all believers value the truth; some don't. The claim that *belief* aims at the truth does not entail that all believers value or aim at the truth; e.g., I may not endorse my own tendency to hold beliefs.²⁹) Because knowledge plausibly requires true belief, we can conclude that likewise to be a successful knower is, among other things, to represent the world accurately. It thus appears that we can discover some constitutive epistemic values without reflecting on the social or political context of knowledge. Why not just continue in this fashion to provide the desired necessary and sufficient conditions for knowledge?

It may be that some constitutive epistemic values (such as truth) can be discovered without a consideration of contextual values, while others require attention to social context. But before granting that even this small part of epistemology can proceed without attention to social and political matters, I think it is valuable to reconsider the basis for regarding truth as an epistemic value. After all, truth may be a constitutive goal of belief, but is there some reason we should see ourselves as committed to forming *beliefs* (as opposed to, say, *acceptings*)? Is there some value in being a *believer*? Should we push the normative project back one more step?

VIII. Cognitive Values for Beings Like Us

In the discussion thus far I have suggested that we should approach the question: What are the truth conditions for knowledge claims? by looking first at our epistemic practices to determine what we do with knowledge attributions, and what legitimate purpose might be served by them. At least in many contexts we use epistemic attributions and judgments in ways that are evaluative: in saying that *S* knows that *p*, I am saying that *S* has met certain cognitive/doxastic standards, with respect to *p*, where those standards capture something of cognitive value. But it isn't clear what the standards should be until we have a clear idea of the value or values at stake. So my suggestion is that we should begin our investigation of what it is to know by investigating what it is that is cognitively valuable (for creatures like us?), i.e., what kinds of cognitive processes, states, and activities, we should endorse. But how do we decide what is cognitively valuable for creatures like us? And should everything that has cognitive value—even some-

thing that has only instrumental value—be properly considered relevant to an analysis of knowledge?

I'd like to be pluralistic about value in general and cognitive value more specifically. But pluralism comes in many forms, and to assert that we should be context-sensitive in judging what is valuable, or that different sorts of things might be instrumentally valuable depending on background goals or purposes, is compatible with claiming that certain things, in fact many things, are intrinsically valuable. For example, I am inclined to say that having beliefs is a good thing for human beings for many reasons, but one reason is that beliefs are an essential component of a kind of agency that is intrinsically valuable.³⁰

I'm not going to offer an argument to defend my suggestion that it is a kind of (moral/autonomous) agency that is the over-arching value. I'm not sure I even have an argument for it yet. But the ideal of agency can at least function as a place-holder in my discussion, for my main goal is to sketch a broad structure for an account of epistemic value generally and knowledge more specifically. To emphasize the structural point, let me contrast my view with two others. The first is the view that epistemic/cognitive virtues constitute a realm of *sui generis* epistemic value, e.g., that truth is good simply for its own sake. The second is the view that epistemic virtues are instrumentally valuable, e.g., we should value truth and so belief (merely?) for survival purposes (the idea being that we need to know what the world is like so we can get around in it). Viewing truth as good because it is conducive to survival strikes me as insufficient, for empirical adequacy, or not too grossly false beliefs, are enough for survival. And although I don't have an argument against a realm of *sui generis* epistemic value, it is not a realm that I find compelling. Instead, I'd like to claim that something is epistemically valuable if it is a cognitive disposition, ability, or achievement that figures in a kind of (moral, autonomous) agency that is intrinsically good.³¹ An analogy here may be helpful, for the notion of a complex intrinsic good that I have in mind is somewhat Aristotelian. *Eudaimonia* is intrinsically good; and yet it is a complex good; other things are intrinsically good because they are constituent parts of a *eudaimon* life. These further goods are not merely instrumental goods, even though their goodness is conditional, i.e., they are good *conditional* on the kind of being we are, but they are *intrinsically* good for us.

Having beliefs is intrinsically good because it is a constitutive requirement of a kind of agency that is part of a *eudaimon* life; because truth is a constitutive value of belief, then we ought also to value truth. Of the many sorts of epistemic evaluations we make, our knowledge attributions seem to be specifically oriented to evaluation of belief: in saying that S knows that p, I am saying that S has met certain cognitive/doxastic standards, with respect to the *belief* that p; because truth is a constitutive goal of belief, this partly explains why we take truth to be a primary condition for knowledge. But what else should we value in the cognitive domain, and more specifically in the domain of belief? After all, it is normally thought that knowledge is not simply true belief. Is there something more? Traditionally, of course, the "something more" beyond true belief has been jus-

tification. And controversies have raged over the nature of justification. Pursuing the analytic strategy I outlined above, the prior question should be: What is at stake in seeking justification? Why are we interested in justified true belief rather than simply true belief?³²

There is a tradition in epistemology which understands justification to be a matter of fulfilling epistemic responsibility. One is justified in believing *p* iff one has been epistemically responsible in forming and maintaining the belief that *p*. We might ask: What is the point of pursuing epistemic responsibility? We should also ask: What counts as epistemic responsibility? What one requires of epistemic responsibility may depend a lot on what view one has about the self and the self's relation to others. E.g., one may think that epistemic responsibility requires that one undertake a solipsistic foundational justification for *p* because doing anything less (such as relying on the reports of others) is irresponsible. Or one may think that epistemic responsibility requires that one consult with others about whether they also believe that *p* (because one's own access to truth is presumed limited). Even the issue between internalists and externalists may arise here: an internalist will want the conditions for epistemic responsibility to be ones that we can self-monitor because we shouldn't be held epistemically responsible for what we have little or no control over, and the externalist will not care about self-monitoring because (to put it crudely) value is placed more on effective agency rather than autonomous agency.

There is a lot of room here for debate over how we should think of ourselves and our cognitive situation, and what we should value cognitively. One might argue that all that matters is truth (though I think they'd have to add, and the avoidance of error), but there's still the question why truth matters. Truth matters *for beings like us* (and not for lots of other sorts of beings), because we have certain capacities (for representing the world and acting on our representations), and the exercise of these capacities is intrinsically good. And although I haven't offered here a view about why justification matters, the strategy I've been pursuing would have us look to the role of justification in informed and autonomous agency.

Note, however, that my intention is not to suggest that what we really need is a "virtue epistemology" (or that we don't); the competing ideas of epistemic responsibility and the conditions for justification they generate may be cashed out in terms that do not bring in virtues.³³ At least it is no part of my view that the primary locus of epistemic evaluation ought to be persons or dispositions rather than particular cognitive states such as beliefs; though I do think that we must consider the fact that the cognitive states are states of beings of a certain sort in order to properly evaluate them. (In other words, my reference to Aristotle is intended to highlight the notion of a conditional intrinsic good, not to invoke his entire ethical legacy.)

So what is my point? How is feminist inquiry into gender and the sexism of our epistemic practices relevant to this project? My suggestion is that questions of value are already implicit in traditional epistemological debates, and that these

questions should be raised more explicitly. And feminist work on the self, on agency, and on social/political values can fruitfully inform and engage these debates. Can feminist inquiry tell us something about what conditions we should include in an "analysis" of *S* knows that *p*? If you think that the question is "what is our concept of knowledge?" you might want to do the traditional apriori investigation and spend a lot of time thinking about fake barns; if you think that the question is "what is the natural kind underlying our epistemic evaluations?" you might want to map out the core cases and then do psychology. But if you aren't convinced that what we do value epistemically is what we ought to value, and if you think (as I do) that there is no reason to favor a natural kind over a non-natural kind as the basis for our evaluations (e.g., the supervenience base for our evaluations may be highly disjunctive), then the question is: what should our concept of knowledge be? But once this is the question, then an adequate definition of knowledge will depend on an account of what is cognitively valuable for beings like us, which raises moral and political issues on which feminists have much to contribute.

One might complain at this stage that my discussion has been too schematic to be convincing. Couldn't one easily maintain that although discussion of what's cognitively valuable has its place, the point and value of knowledge is best captured by regarding it simply as a matter of reliably formed true belief? Then my suggestion that feminist inquiry into our social and cognitive lives is essential to normative epistemology has little or no bite.

This is not the place to debate the virtues of reliabilism. My argument is intentionally schematic for my point is not to settle any debates but to open up space for further debate. Methodologically my response to the objection is that a defense of reliabilism cannot simply assume that reliably formed true belief constitutes the point and value of knowledge and would need to address the social and moral dimension of belief. For example, consider that although a reliabilist account of knowledge captures some of our concern with knowledge, it has often been observed that reliability is not all that matters in forming beliefs—if it were, then it would make epistemic sense to adopt a method of believing only tautologies. Falsehood is certainly an epistemic vice, but if avoiding falsehood were the single overriding epistemic virtue, then it could make sense to believe nothing.³⁴ But this makes no sense if our concern is what's cognitively valuable for beings like us, and if we also attach significant value to (informed) agency, for perfectly reliable processes of belief formation provide agents no meaningful basis for action (one cannot act on tautologies alone!).

Moreover, when asked why reliability is cognitively valuable, a plausible answer will point to the value of reliable information in enabling us to act. But this answer is weak if the connection between true belief and action is purely instrumental; as mentioned before, reliably true beliefs are not necessary for getting around effectively. Rather there is a deeper connection between our cognitive lives and our practical lives. Reliable methods are valuable because they produce true beliefs, and true beliefs are necessary to achieve a kind of agency

that is good; likewise having a coherent and fruitful system of beliefs is good because it promotes the effective exercise of a kind of agency that is good. An adequate account of knowledge must take a variety of characteristics into account, e.g., truth, coherence, reliability, informativeness, fruitfulness, etc. Again, the point here is not to suggest that reliable and fruitful methods, coherent belief sets, etc., are instrumentally good because they serve the political goal of furthering agency. Rather, the point is that they are intrinsically good for beings like us; their intrinsic goodness is conditional on certain facts about us as moral agents.

But once we begin to think in a more robust way about the value of agency and the cognitive life of a flourishing agent, it becomes clear that there are also more general issues at stake: how should we organize ourselves and our cognitive activities within communities so as to promote effective and informed agency? What is cognitively valuable for us *as a group*, i.e., how might we best cooperate in our cognitive efforts, if we value the capacity in each individual to exercise their agency? (Or to achieve *eudaimonia*, or some other intrinsic good?) For example, there has been considerable discussion in the context of feminist moral theory of the notion of autonomy. The charge has been that traditional moral theory has not been sufficiently attentive to the social requirements for and limitations on autonomy. The charge is that not only have certain (exaggerated) ideals of independence and self-sufficiency been overrated, but valuable and sometimes unavoidable forms of interdependence have been ignored and/or scorned. A parallel discussion has emerged in the context of feminist epistemology: the lone epistemic agent is in some important sense a myth (wolf children are not plausible models for moral or epistemic life). Not only are we dependent on others for what we know, but our epistemic interdependence is a good thing; but at the same time we should be attentive to the value of epistemic autonomy.³⁵

At this stage of my discussion I don't want to get involved in the details of these debates, but I think feminist theorizing has been effective in showing that different sets of norms and practices "construct" different kinds of knowers. An important task for epistemology, as I see it, is to consider the full range of norms, practices, and conventions, that together enable autonomous cognitive/epistemic agency. So, e.g., even if we allow that there is a refined conception of epistemic independence that is valuable—one that recommends having a "mind of one's own", taking responsibility for one's own beliefs, etc.—and that this is a factor in achieving epistemic autonomy, epistemic autonomy is not just a matter of being independent in this sense, for it requires participation in a social network with other epistemic agents, whose own agency (and independence) provides a crucial check on our own beliefs.³⁶

Here one might again object that I have managed to show how feminist work is relevant to the task of defining knowledge only by conflating that task with the project of describing the norms and practices that, in a particular context, enable one to be a more successful knower. E.g., a reliabilist could grant that it is both difficult and important to determine what epistemic practices—considered both individually and socially—are more likely to enable individuals to gain knowl-

edge. Because our epistemic practices are very messy affairs that vary tremendously across context, feminist inquiry may help us uncover some of the limitations of our actual practices and propose more effective ones. But none of this requires us to incorporate feminist insights into effective epistemic norms into our *analysis* of knowledge. The proper analysis of knowledge should simply focus on truth-tracking; the rest—including how we think of ourselves and our relations to each other—is just heuristics.³⁷

I'm happy to grant that on a reliabilist view, much of what I am counting as relevant to determining the proper analysis of knowledge falls under the rubric of heuristics. But to consider the issue from the point of view of a reliabilist is to beg the question—for what recommends reliabilism over the various alternative epistemic positions? On what basis do we value truth, and grant it place as the dominant value in our cognitive practices? Even if in the end reliabilism turns out to be the most compelling epistemic view, some explanation must be given of its normative grip on us, of why we ought to evaluate ourselves and others in its terms.

I've suggested that one way of thinking about epistemic value is to resist the suggestion that epistemic values need be either instrumental or *sui generis*; there may be a more basic value of which they are constituent parts. Although in this paper I've pointed to the value of a certain kind of agency as a basis for adjudicating what's epistemically valuable, I'm far from certain this can be spelled out, and in fact I am tempted by other alternative framings of our epistemic evaluations.

My remarks here are not sufficient to provide a very clear example of what sorts of conditions on knowledge we might consider, informed by the feminist research I've been alluding to. I'm afraid I haven't gotten quite that far in the program! My own inclination would be to resist the suggestion that we need to focus on one overall evaluative notion in this area: knowledge. This allows me to bring a pragmatist theme in my discussion back to the surface: on a critical analytical approach to the question, "What is knowledge?" we begin by asking what we need the concept for, what work it is doing for us. I've argued that the theoretical work we want it to do is primarily evaluative: to say that S knows that p is to offer an evaluation of S's belief that p with respect to standards for belief formation and retention. This allows that there may be other jobs that need to be done, and some may want to retain the term 'knowledge' for those jobs; or one may prefer to reserve the term 'knowledge' for our current concept. That's fine with me; I don't want to quarrel about who gets to use the term. But even considering the evaluative work that needs to be done, there are a variety of different purposes one might have in making epistemic evaluations. I've suggested that one way of thinking about epistemic value is to understand it in the context of a broader notion of autonomous agency (though I've only briefly sketched how this might work). But here too, I want to allow that there are many different kinds of value and that alternative conceptions of epistemic value may depend on those. Ultimately I would hope for a proliferation of epistemic notions. Again, I'm not intending to close off inquiry, but to open it up.

However, the normative question, I think, is not in the end optional: a discussion of the truth conditions for knowledge claims that does not critically reflect on the broader purposes for our use of the concept, and that does not take up the issue of epistemic value is impoverished. But once we do engage in this critical reflection, feminist research into the sexism in our current practices and into alternative conceptions of agency and value, become highly relevant and important.

Notes

1. Thanks to Elizabeth Anderson, Donald Baxter, Cheshire Calhoun, Margaret Gilbert, Jim Joyce, Jeffrey Kasser, Krista Lawlor, Ruth Millikan, Peter Railton, Laura Schroeter, David Velleman, and Stephen Yablo for helpful discussion of issues taken up in this paper. Special thanks to Karen Jones for excellent commentary on an earlier version presented at the Central Division meetings of the American Philosophical Association, May, 1998. And thanks to audiences at the University of Connecticut and Colby College for their helpful discussion. An earlier version of this paper also appears as, "Defining Knowledge: Feminist Values and Normative Epistemology," in the Proceedings of the World Congress of Philosophy, 1998.
2. One of the "notable exceptions," is Helen Longino, "The Fate of Knowledge in Social Theories of Science," in F. Schmitt, ed., *Socializing Epistemology: The Social Dimensions of Knowledge* (Lanham, Maryland: Rowman and Littlefield, 1994), pp. 135–157.
3. Allowing for differences in conceptual resources across time and place, it might be important to specify the group whose concept is at issue: what is the dominant concept of knowledge within such and such a culture at such and such a time? However, it is notable that some traditional epistemologists seem to assume that there is *one* true concept of knowledge and the question is what is the proper analysis of that.
4. See, e.g., Jaegwon Kim, "What is "Naturalized Epistemology"?" *Philosophical Perspectives* 2 (1988): 381–405.
5. This is a slight revision of the slogan Stephen Stich uses to characterize "analytic epistemology." See Stephen Stich, *The Fragmentation of Reason* (Cambridge, MA: MIT Press, 1993), p. 93.
6. A valuable summary of the kind of research I have in mind can be found in Elizabeth Anderson, "Feminist Epistemology: An Interpretation and a Defense," *Hypatia* 10:3 (Summer 1995): 50–84. A full bibliography of recent work in feminist epistemology would no doubt be helpful, but not something I can provide here. I encourage the reader to consult the bibliography of Anderson's paper. In addition to other feminist work I've cited in this essay, there are three recent collections that are especially useful: L. Antony and C. Witt, ed., *A Mind of One's Own* (Boulder, CO: Westview, 1993); L. Alcoff and E. Potter, ed., *Feminist Epistemologies* (New York: Routledge, 1993); and the special issue of *Philosophical Topics* 23:2 (Fall 1995) on "Feminist Perspectives on Language, Knowledge, and Reality." These collections provide exposure to a broad range of authors and through their bibliographies to the extensive work in feminist epistemology.
7. Louise Antony has made this point in several contexts, e.g., in "Comment on Naomi Scheman," *Metaphilosophy* 26, no. 3 (July 1995): 191–198.
8. Note that Stich considers this move in response to the fact of "cognitive diversity" in *The Fragmentation of Reason*, pp. 80–82.

9. See also Barry Stroud, *The Significance of Philosophical Skepticism* (Oxford: Oxford University Press, 1984), pp. 57ff; also J. Joyce "The Lasting Lesson of Skepticism," (1998 manuscript).
10. J.L. Austin, "Other Minds," in *Philosophical Papers*, second edition, ed., J.O. Urmson and G.J. Warnock (Oxford: Oxford University Press, 1970), p. 99.
11. "Other Minds," p. 100.
12. Although Austin himself focused on first person knowledge claims, his suggestions might be extended to second and third person attributions as well. I don't mean to suggest that an expressivist account of knowledge would be easy to provide, for all the reasons expressivist accounts are difficult (e.g., how to deal with embedding, connectives, etc.). Nor am I convinced that Austin's account is best understood as expressivist. Rather, I'm suggesting that in asserting that S knows that p one may also be performing another illocutionary act concerned with authorization.
13. *Fragmentation of Reason*, Ch. 4.
14. *Fragmentation of Reason*, pp. 87–89.
15. I take it to be a broad assumption in philosophical approaches to epistemology that the goal is to provide a normative account. (E.g., see Kim, "What is "Naturalized Epistemology"?) Though some see there to be two projects, one non-normative and another normative, e.g., Alvin Goldman, "Epistemic Folkways and Scientific Epistemology," in *Liaisons: Philosophy Meets the Cognitive and Social Sciences* (Cambridge, Mass: MIT Press, 1992), pp. 155–75, reprinted in Paul K. Moser and Arnold vander Nat, ed., *Human Knowledge*, second edition (New York and Oxford: Oxford University Press, 1995), pp. 439–453. Note, however, that those who resist a normative inquiry are not in a good position to complain that feminist discussions are "just sociology, not philosophy," for the same question might be asked of their work: what makes it philosophy rather than sociology or psychology. In short, if feminists are acknowledged to be making useful sociological/psychological observations, then the non-normative epistemological project should take those observations seriously; but (as I will argue) if the project is normative, then there too feminists have something to contribute to the discussion of value.
16. Can a very sophisticated but still narrow "reflective equilibrium" method capture the normative dimension here, and so do justice to the question of epistemic value? Stich argues not, and I would agree, since in principle what matters is not that we achieve a "reflective equilibrium" amongst our ordinary epistemic intuitions/judgments but that we determine what's valuable, and what's valuable may not show itself in our epistemic intuitions/judgments. Things will have to get more complicated to spell out this argument, however, since what we ought to value may be something that we can only determine via a method of wide reflective equilibrium. I will discuss this further below. My point here need only be that a narrow reflective equilibrium that takes into account only our intuitions/judgments regarding epistemic matters is insufficient.
17. On the idea of reason as reflective self-government, see e.g., Elizabeth Anderson, "Feminist Epistemology," pp. 52–3. Lest one be concerned that the introduction of value into epistemology render it non-objective, consider that many hold that there are standards for objective normative inquiry. See also Elizabeth Anderson, "Knowledge, Human Interests, and Objectivity in Feminist Epistemology," *Philosophical Topics* 23:2 (Fall 1995), esp. pp. 32–37.
18. I use the term "analytical" for this approach because of its use in contemporary feminist theory to designate such a project. In particular I have in mind Joan Scott's important essay, "Gender: A Useful Category of Historical Analysis," *American His-*

- torical Review 91:5 (1986) 1053–75. Sandra Harding also employs this term, e.g., in “The Instabilities of the Analytical Categories of Feminist Theory,” in Sandra Harding and Jean F. O’Barr, eds., *Sex and Scientific Inquiry* (Chicago and London: Univ. of Chicago Press, 1987), pp. 283–302, and Sandra Harding, *The Science Question in Feminism* (Ithaca and London: Cornell Univ. Press, 1986), e.g., p. 30. Note that my use of the term differs from Stich’s in his phrase, “analytic epistemology.”
19. My own temptation to view philosophical questions “analytically” is partly due to the influence of Paul Grice. Consider, e.g., “Reply to Richards,” in R. Grandy and R. Warner, ed., *Philosophical Grounds of Rationality* (Oxford: Oxford University Press, 1986), pp. 45–106, and “Method in Philosophical Psychology (From the Banal to the Bizarre),” *Proceedings of the American Philosophical Association*, vol. 78 (November 1975): 23–53. Grice provides one of the important recent philosophical precedents for this approach.
 20. In the limiting case of the analytic project, the meaning of the term is simply stipulated without reference to standard examples—this happens when a new theoretical term or notation is introduced—though in many projects there is some concern to draw on and address common (ordinary or theoretical) usage. In the latter case, it is within the scope of such a project both to challenge the idea that we need any such concept i.e., that there is valuable work to be done in the vicinity, and to challenge the particular concept that has been employed to do it.
 21. On different sorts of definition and the possibility or impossibility of revisionary analytical projects, see W. V. O. Quine, “Two Dogmas of Empiricism,” in *From a Logical Point of View* (New York: Harper and Row, 1963) section 2, and R. Rorty, “Metaphilosophical Difficulties of Linguistic Philosophy,” introduction to *The Linguistic Turn* (Chicago and London: University of Chicago Press, 1992) pp. 1–39.
 22. Just as descriptive and conceptual projects cannot be kept entirely separate, analytic projects cannot be kept entirely separate from either. For example, descriptive projects unearth important and often novel similarities and differences in the items under consideration, and often provide innovative explanations aiming to do justice to the richness of the phenomena. But there comes a point when it’s necessary to re-evaluate the conceptual tools the project started with, and to focus the inquiry by reassessing its aims; this invites a shift to a more analytical mode. Analytical projects reflect on, evaluate, and revise the conceptual tools we have for organizing phenomena, but in order to assess realistically what tools we need, and why we need them, they depend crucially on descriptive efforts. Stated crudely, a descriptive project begins with a rough conceptual framework in place, and looks to the world to fill in the details; an analytic project begins with a rough understanding of the salient facts, and works to construct a conceptual framework that can offer a useful way of organizing them. In both cases we do well to have a sensitive and thorough understanding of our existing concepts.
 23. Consider, e.g., Miranda Fricker, “Rational Authority and Social Power: Towards a Truly Social Epistemology,” *Proceedings of the Aristotelian Society*, vol., XCVII, Part 2, (1998): 159–177; and Edward Craig, *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* (Oxford: Oxford University Press, 1990).
 24. Helen Longino, *Science as Social Knowledge* (Princeton: Princeton University Press, 1990), pp. 4–7.
 25. Consider, e.g., R. Rorty, “Solidarity or Objectivity?” in *Objectivity, Relativism, and Truth: Philosophical Papers*, vol. 1 (Cambridge U. Press, 1991), pp. 21–34.
 26. For a useful discussion of truth-directedness as an “internal” goal of belief, see P. Railton, “Truth, Reason, and the Regulation of Belief,” *Philosophical Issues* 5 (1994):

- 72–75; and David Velleman, "The Possibility of Practical Reason," *Ethics* 106 (July 1996): 694–726, esp. pp. 708–714.
27. See "Truth, Reason, and the Regulation of Belief," p. 74.
 28. For valuable discussion and elaboration of this point see David Velleman, "How Belief Aims at the Truth," (manuscript, May 1998).
 29. Again, Velleman's work is very helpful in explaining this point. See "The Possibility of Practical Reason," pp. 709–10.
 30. Thanks to David Velleman for conversation on this and related points. For further discussion of links between epistemic concerns and agency, especially in the skeptical tradition, see Christopher Hookway, *Scepticism* (Routledge 1990), pp. 132–136. Thanks to Jeffrey Kasser for bringing Hookway's book to my attention.
 31. Note that to flesh out this picture I'd have to make a case that the kind of agency that is intrinsically valuable requires belief/truth rather than acceptance/empirical adequacy. This task goes well beyond the scope of this paper.
 32. William Alston's paper, "Epistemic Desiderata," *Philosophy and Phenomenological Research*, vol. 52, no. 3 (September 1993): 527–551, is very useful in challenging the idea that there is an ordinary conception of justification that our epistemic inquiry must do justice to, and discusses several different purposes for the notion. Alston's argument is effective in showing that participants in the contemporary debate over justification may well be talking past each other. I'm sympathetic to Alston's methodological project, though my turn to a broader conception of value to resolve some of the questions is one he doesn't consider.
 33. See, e.g., Linda T. Zagzebski, *Virtues of the Mind: An Inquiry into the Nature of Virtue and The Ethical Foundations of Knowledge* (Cambridge: Cambridge University Press, 1996). Thanks to Karen Jones for bringing Zagzebski's work to my attention.
 34. These points are eloquently put by W. James in his "The Will to Believe," in *The Will to Believe and Other Essays* (New York: Dover Publications, 1956). See also Hartry Field, "Realism and Relativism," *Journal of Philosophy* (1982), p. 565–6.
 35. Louise Antony has been good at reminding us all of this. See, e.g., Louise Antony, "Sisters Please, I'd Rather Do It Myself: A Defense of Individualism in Feminist Epistemology," *Philosophical Topics* 23, no. 2 (Fall 1995): 59–94.
 36. Thanks to Peter Railton for discussion of this issue.
 37. Thanks to Karen Jones for raising this objection.

References

- Alcoff, Linda and Elizabeth Potter, eds., *Feminist Epistemologies* (New York: Routledge, 1993)
- Alston, William P. "Epistemic Desiderata," *Philosophy and Phenomenological Research* 52: 3 (September 1993): 527–551.
- Anderson, Elizabeth. "Feminist Epistemology: An Interpretation and a Defense," *Hypatia* 10:3 (Summer 1995): 50–84.
- . "Knowledge, Human Interests, and Objectivity in Feminist Epistemology," *Philosophical Topics* 23:2 (Fall 1995): 27–58.
- Antony, Louise. "Comment on Naomi Scheman," *Metaphilosophy* 26, no. 3 (July 1995): 191–198.
- . "Sister's Please, I'd Rather Do It Myself: A Defense of Individualism in Feminist Epistemology," *Philosophical Topics* 23:2 (Fall 1995): 59–94.
- and Charlotte Witt, eds., *A Mind of One's Own* (Boulder, CO: Westview, 1993).
- Austin, John L. "Other Minds," in *Philosophical Papers*, second edition, ed., J.O. Urmson and G.J. Warnock (Oxford: Oxford University Press, 1970), pp. 76–116.

- Craig, Edward. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* (Oxford: Oxford University Press, 1990).
- Field, Hartry. "Realism and Relativism," *Journal of Philosophy* 79 (1982): 565–6.
- Fricker, Miranda. "Rational Authority and Social Power: Towards a Truly Social Epistemology," *Proceedings of the Aristotelian Society* 97, Part 2, (1998): 159–177.
- Goldman, Alvin. "Epistemic Folkways and Scientific Epistemology," in *Human Knowledge*, second edition, ed., Paul K. Moser and Arnold vander Nat (New York and Oxford: Oxford University Press, 1995), pp. 439–453, reprinted from *Liasons: Philosophy Meets the Cognitive and Social Sciences* (Cambridge, Mass: MIT Press, 1992), pp. 155–75.
- Grice, H. Paul. "Reply to Richards," in *Philosophical Grounds of Rationality*, ed., R. Grandy and R. Warner (Oxford: Oxford University Press, 1986), pp. 45–106.
- . "Method in Philosophical Psychology (From the Banal to the Bizarre)," *Proceedings of the American Philosophical Association* 78 (November 1975): 23–53.
- Harding, Sandra. "The Instabilities of the Analytical Categories of Feminist Theory," in *Sex and Scientific Inquiry*, ed., Sandra Harding and Jean F. O'Barr (Chicago and London: Univ. of Chicago Press, 1987), pp. 283–302.
- . *The Science Question in Feminism*, (Ithaca and London: Cornell Univ. Press, 1986).
- Haslanger, Sally, ed., *Feminist Perspectives on Language, Knowledge, and Reality*, special issue of *Philosophical Topics* 23:2 (Fall 1995).
- James, William. "The Will to Believe," in *The Will to Believe and Other Essays* (New York: Dover Publications, 1956).
- Joyce, James. "The Lasting Lesson of Skepticism," (1998 manuscript).
- Kim, Jaegwon. "What is "Naturalized Epistemology"?" *Philosophical Perspectives* 2 (1988): 381–405.
- Longino, Helen. "The Fate of Knowledge in Social Theories of Science," in *Socializing Epistemology: The Social Dimensionso of Knowledge*, ed., F. Schmitt (Lanham, Maryland: Rowman and Littlefield, 1994), pp. 135–157.
- Quine, W. V. O. "Two Dogmas of Empiricism," in *From a Logical Point of View* (New York: Harper and Row, 1963), pp. 20–46.
- Railton, Peter. "Truth, Reason, and the Regulation of Belief," *Philosophical Issues* 5 (1994): 72–75.
- Rorty, Richard. "Metaphilosophical Difficulties of Linguistic Philosophy," in *The Linguistic Turn* (Chicago and London: University of Chicago Press, 1992), pp. 1–39.
- . "Solidarity or Objectivity?" in *Objectivity, Relativism, and Truth: Philosophical Papers, vol. 1*, (Cambridge U. Press, 1991), pp. 21–34.
- Scott, Joan. "Gender: A Useful Category of Historical Analysis," *American Historical Review* 91:5 (1986) 1053–75.
- Stephen Stich, *The Fragmentation of Reason* (Cambridge, MA: MIT Press, 1993).
- Stroud, Barry. *The Significance of Philosophical Skepticism* (Oxford: Oxford University Press, 1984).
- Velleman, David. "The Possibility of Practical Reason," *Ethics* 106 (July 1996): 694–726.
- . "How Belief Aims at the Truth," (manuscript, May 1998).
- Zagzebski, Linda T. *Virtues of the Mind: An Inquiry into the Nature of Virtue and The Ethical Foundations of Knowledge* (Cambridge: Cambridge University Press, 1996).