

EPISTEMIC EVALUATION AND THE AIM OF BELIEF

Kate Nolfi

A thesis submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Master of Arts in the Department of Philosophy.

Chapel Hill
2010

Approved by:

Ram Neta

Matthew Kotzen

William Lycan

© 2010
Kate Nolfi
ALL RIGHTS RESERVED

ABSTRACT

Kate Nolfi: Epistemic Evaluation and the Aim of Belief
(Under the direction of Ram Neta)

In this paper, I articulate a popular and, I think, promising methodological approach to developing an account of epistemic evaluation. I then sketch the account of epistemic evaluation that I argue results from properly implementing this methodological approach. In so doing, I argue that if one accepts the methodological approach I articulate here, then one should also accept that the constitutive aim of belief is not, as many philosophers have assumed, appropriately cashed out in terms of truth. Rather, one should endorse the claim that beliefs aim at being well-suited to serve in practical reasoning that yields successful action. Finally, I briefly address two lines of resistance to the pragmatically driven characterization of the constitutive aim of belief and the resulting account of epistemic evaluation that I develop here.

ACKNOWLEDGMENTS

Thanks to Ram Neta for all his work in advising this project. Whatever insight I have achieved here would certainly have been impossible without the help of Ram's careful thinking and critical engagement. I am also grateful to Bill Lycan and especially to Matt Kotzen for both of their invaluable assistance in developing and refining the ideas presented here. Finally, I would like to thank my peers in the Philosophy Department at UNC for their unbounded intellectual support.

TABLE OF CONTENTS

Section	Page
Introduction.....	1
1. Articulating a Methodological Approach.....	2
1.1.....	2
1.2.....	4
1.3.....	8
1.4.....	10
2. Implementing the Proposed Methodology: An Account of Epistemic Evaluation.	12
2.1.....	13
2.2.....	16
2.3.....	18
2.4.....	22
3. Considering Objections	25
3.1.....	25
3.2.....	30
Conclusion	39
References.....	41

Introduction

In this paper, I aim to articulate a methodological approach to epistemological inquiry that I take to be relatively popular, at least in its broadest formulation, among contemporary epistemologists. I rely heavily on the work of Ernest Sosa in articulating this methodological approach as Sosa's work explicitly and transparently (relatively speaking) attempts to implement the methodology on offer here. I will not, however, offer a defense of this methodological approach in this paper. Instead, I go on to develop an account of epistemic evaluation by carefully implementing the methodology articulated here. However, the account of epistemic evaluation that I argue results from carefully implementing the methodological approach on offer here characterizes epistemic evaluation in pragmatic terms rather than in terms of a truth-goal. Finally, I defend this account against two related lines of objection that the epistemologist sympathetic to an account cashing out epistemic evaluation in terms of truth might level against the pragmatically-driven account I develop below. In so doing, I hope to show that the epistemologist who endorses the methodology articulated here should prefer the pragmatically-driven account of epistemic evaluation that I develop in what follows over a more traditional account cashing out epistemic evaluation in terms of a truth-goal. So, if one accepts this methodological approach to epistemological inquiry, then one should also accept the account of epistemic evaluation that I advance here.

1. Articulating a Methodological Approach

1.1

The methodological approach that I pursue in what follows takes as its starting point that beliefs are appropriately understood as a kind of performance. As a result, the methodology on offer here recommends that the epistemologist ground her account of specifically epistemic evaluation—evaluation of belief—in an account of performances and performance evaluation given in general terms. In this spirit, I aim here to lay the foundation for an account of epistemic evaluation by articulating Sosa's account of performances and performance evaluation.

Sosa takes the performance that an archer gives in shooting an arrow at a target as a paradigm case in his own discussions. However, it is important to note that Sosa understands the notion of a performance quite broadly.¹ Thus, the heart, in pumping blood, is engaged in a kind of performance and, similarly, for Sosa, beliefs are

¹ It may be helpful here to explicitly distinguish two types of performers who give or engage in performances: organisms (e.g. the archer, the dancer) and systems (e.g. the circulatory system, the cognitive system) operating within an organism. It may be right to say that a performance given by a performer of the first sort inherits its constitutive aim from the performer's aim in giving the performance. However, I suggest that, in the case where the performer is of the second sort, the aim of the performer's particular performance is derived from the purpose or function of the performing system operating within the organism.

performances of the cognitive system.² Specifically, belief or believing is a performance that is the product of the operation of a cognitive system within an organism and so what it is to be a belief is just, on Sosa's view, to be a performance of the cognitive system of some organism.

Characterized in the most basic terms, performances are actions or behaviors that aim to achieve some purpose or end. Any particular performance is the kind of performance that it is in virtue of having the particular aim that it does. The archer's shot, for example, aims to hit the target. The heart's beating aims, let's say, to circulate blood.³ Moreover, the archer's shot counts as a *shot* (rather than a mere release of the arrow from the bow) in virtue of its being aimed at hitting the target. Similarly, part of what it is to be heartbeat is to be a performance aimed at circulating blood through an organism.⁴ The aim of a performance is at least partially constitutive of that performance's being the kind of performance that it is. Put in other terms, a performance is not just any doing, but a doing in order to or for the purpose of Φ .

Often, we can discover the constitutive aim, end, or purpose of a particular type of performance by coming to understand when and why performers who give that kind of performance do so. We discover that an archer's shot aims to hit the target by

² Thus, although his examples (e.g. the archer, the dancer) do not always make this point clear, taking belief to be a kind of performance does not commit the epistemologist endorsing this methodology to epistemic voluntarism.

³ It is worth pointing out that the aim of the heart's beating is probably a great deal more complex than I have suggested here.

⁴ Organ systems, I take it, are identified by their proper function and so an organ system cannot have multiple functions (although, of course, two organ systems—each characterized in functional terms—might well be realized in some of the same physical structures). As a result, performances of an organ system will have a single, constitutive aim, which derives from the proper function of the performing system.

discovering when and why archers shoot arrows. Similarly, we discover that the heart's beating aims to pump blood by discovering when and why organisms have circulatory systems (i.e. by discovering the way in which the circulatory system serves to maintain or increase the biological fitness of the organism within which it operates). And so, because she understands beliefs as performances of the cognitive system, the epistemologist adopting the methodology I am developing here maintains that we discover the constitutive aim of belief by discovering the way in which the cognitive system serves to increase or maintain the biological fitness of the organism within which it operates.

1.2

The account of performance evaluation that I adopt in what follows maintains that, since performances are identified as the particular kind of performance that they are by their constitutive aims, performances are always appropriately evaluated in reference to their aims. In an effort to flesh out this thesis, Sosa proposes that performances have AAA normative structure. That is, performances are appropriately evaluated along three dimensions: accuracy, adroitness, and aptness.⁵ Most broadly construed, the degree to which a performance is accurate is simply the degree to which the performance successfully meets its aim. The degree to which a performance is adroit is just the degree to which the performance manifests skill on the part of the performer. Finally, the degree to which a performance is apt is the degree to which the performance is accurate *because* it is adroit. That is, a performance is apt just in case the performance meets its aim because it manifests the performer's skill and so if a performance is apt, then it is

⁵ See Sosa, E. (2007), Lecture 2: A Virtue Epistemology, pp. 22-45.

appropriate to attribute the success of the performance to the performer since the performer's skill causes the performance to be successful.

Consider Sosa's archer.⁶ The archer's performance is accurate if he hits his target, his performance is adroit if it manifests the archer's skill in executing the shot, and the performance is apt just in case the archer's arrow hits the target *because* of his skill in executing the shot. If the archer aims and then releases the arrow from the bow in a skillful manner (so that, under normal conditions, the arrow would have hit the target), but an unexpected gust of wind blows the arrow off course and so the arrow fails to hit the target, then the archer's shot is adroit but not accurate and so not apt. If, on the other hand, a novice aims and releases the arrow without skill, but manages to hit the target by a stroke of luck (let us say that an unexpected gust of wind blows the arrow onto a course that lands the arrow in the center of the target), then the novice's shot is accurate, but not adroit and so not apt. Finally, we might imagine that the archer aims and releases his arrow with skill, an unexpected gust a wind blows the arrow off its original course, but then a second unexpected gust of wind blows the arrow back on course and so the arrow hits the target. In this case, the archer's shot is both accurate and adroit. However, the archer's shot fails to be apt because the archer's skill does not cause (or is not sufficient to cause) the archer's arrow to hit the target—the second unexpected gust of wind plays a crucial role in causing the arrow to hit the target.

Consider now the case where the performer is a system operating within a larger organism. In this sort of case, it is when the system is performing or functioning properly, that its performances succeed in meeting their aim. The circulatory system, for

⁶ See pages 22-23 of Sosa (2007) for Sosa's discussion of the archer.

example, is functioning or performing properly (and so its performances are accurate) to the degree that the heart's beating succeeds in pumping blood.⁷ The heart's beating will be adroit to the degree that the heart's beating occurs in such a way (i.e. with the proper frequency and force) so as to be well suited to circulating blood throughout the body. Finally, the heart's beating will be apt if it succeeds in pumping blood throughout the body as a result of proceeding in a way that is well suited to this task, rather than by some fluke or stroke of luck.

The epistemologist pursuing the approach I have been developing here will apply the account of the accuracy, adroitness, and aptness articulated above to the case of belief and then employ the resulting account of epistemic evaluation in developing accounts of knowledge and justification.⁸ In this vein, the epistemologist endorsing the methodology on offer here maintains that justified belief is just adroit belief. That is, justified belief is belief that results from a cognitive process that reliably, generally, or under normal circumstances produces beliefs that meet their constitutive aim. Sosa, for example, maintains that a belief is justified if and only if it manifests epistemic virtue, and so an account of the epistemic virtues will explain which of our beliefs are justified and why.⁹ Recalling what it takes for a habit to count as an virtue or skill, a subject's belief that P will be justified when the subject's coming or continuing to believe that P is the result of a cognitive habit that, generally speaking or under normal conditions, produces accurate

⁷ It is worth making explicit that this use of the term accuracy is somewhat simulative, and so the reader should not be thrown by the seeming awkwardness of the claim that a performance of the circulatory system is accurate.

⁸ Again, see pages 22-23 of Sosa (2007) or Sosa (2001).

⁹ Therefore, on Sosa's picture, epistemic virtue is explanatorily prior to justified belief and, thus, knowledge.

belief. A belief is accurate to the extent that, as a performance, it meets its aim and so, on most views (Sosa's included) accurate belief is just true belief. So, for Sosa, S is justified in believing that P just when her belief that P results from the operation of a cognitive habit that reliably engenders true belief.

The epistemologist adopting the methodological approach on offer here goes on to propose that knowledge is simply apt belief. So, a subject, S, has animal knowledge that P just in case S believes that P, S's belief that P is accurate (which, recall, on most views is just to say that the belief is true), S's belief that P is adroit, and S's belief that P is accurate *because* it is adroit. The account of epistemic evaluation that results from the application of this methodological approach to epistemological inquiry vindicates the intuition that knowledge requires justification. Since the epistemologist adopting the methodology developed here cashes out justified belief in terms of adroit belief—belief that manifests epistemic virtue or cognitive skill; belief produced by a cognitive process that reliably generates accurate belief—justification, on the resulting picture, is a necessary condition for knowledge. Additionally, by understanding knowledge as apt belief, this account excludes Gettier cases from the domain of knowledge by requiring that the *truth* of S's belief that P *result from so as be explainable by* an exercise of S's epistemic virtue in order for S's belief that P to count as knowledge.¹⁰

¹⁰ In certain texts, when presenting his account of knowledge, Sosa discusses safe belief. On Sosa's view, a belief is safe when it could not easily have been false. More precisely, S's belief that P is safe just in case S would not have believed P if P had not been true. Sosa sometimes seems to suggest that safety is required for knowledge. However, Sosa is perhaps more charitably interpreted as suggesting not that safety itself is a necessary condition for knowledge, but rather that safety is a good indicator of adroitness and so, since animal knowledge require adroit belief, safety is a good indicator of one necessary condition for knowledge. Given Sosa's account of adroit belief as belief that manifests epistemic virtue and Sosa's account of what it takes for a cognitive habit to count as an

On the methodological approach that I have been developing here, epistemology is, most essentially, in the business of supplying an account of epistemic evaluation—an account of how to appropriately evaluate beliefs, understood as the performances of the cognitive system. Since all performances are appropriately evaluated with reference to their constitutive aims, the first task facing the epistemologist adopting the methodology on offer here is the task of identifying the aim of the performances of the cognitive system. Identifying this aim involves identifying the particular job or function that the cognitive system performs in maintaining or enhancing the biological fitness of the organism within which it operates. Thus, this first task is an empirical one: just as the biologist discovers the purpose or function of the circulatory system (composed of the heart, arteries, veins, etc.) by investigating real circulatory systems operating within organisms, so to the epistemologist begins her inquiry by studying cognitive systems operating within organisms in an effort to discover how the cognitive system can and does help to maintain or enhance the biological fitness of the organism.¹¹

epistemic virtue, adroit beliefs will be safe and most safe beliefs will be adroit. However, it is possible for a belief to be safe and true, but fail to be adroit and so fail to meet the criterion for animal knowledge.

¹¹ As Alvin Plantinga has pointed out, the notion of proper function is ambiguous between something like the function that the system was designed to perform and the function that the system currently performs. I adopt the second of these two interpretations here and so maintain that while the function that a system is designed to perform might well provide some evidence for the proper function of the system in question, the proper function of the system need not be identical or even similar to the function that the system was designed to perform. The proper function of a system is given by the particular job or purpose that the system currently performs in order to increase or maintain the biological function of the organism within which it operates.

Armed with suitably detailed characterization of the way in which the performance or functioning of the cognitive system contributes to the biological fitness of the organism within which it operates and, derivatively, a detailed account of the aim of our cognitive performances, the epistemologist can work from the AAA normative structure of performances to develop a suitably detailed account of epistemic evaluation. A performance is appropriately evaluated as accurate just in case that performance meets its aim, and so the epistemologist can say precisely what it takes for a cognitive performance, a belief, to be accurate by referring back to her characterization of the aim of the cognitive system. Similarly, performances are appropriately evaluated as adroit when they result from a habit or process that reliably generates performances that meet their aim. So, the epistemologist develops an account of cognitive skill and, in so doing, an account of epistemic adroitness by, again, looking out into the world to see which of our cognitive processes reliably produce beliefs that meet their performative aim. At this stage as well, then, the epistemologist's project is, at least in part, an empirical one. Nevertheless, in developing an account of cognitive skill, and thus an account of epistemic justification, the epistemologist is engaged in the project of identifying those cognitive processes the performing of which is worth cultivating. The epistemologist adopting the methodology developed here is not engaged in a merely descriptive project. Rather, she is in the business of making recommendations that carry full-blooded normative force with respect to what we should believe and which ways of forming beliefs we should adopt ourselves and try to cultivate in others.

Although I do not aim to defend it here, the methodological approach to epistemological inquiry that I have outlined above has certain appealing features that are worth briefly mentioning. First, in accepting this methodological approach, the epistemologist sets herself a well-defined project and adopts a clear strategy toward this end. The epistemologist aims to give an account of how one can appropriately evaluate beliefs, understood as performances of the cognitive system. Toward this end, the epistemologist must begin by characterizing the constitutive aim or goal of our cognitive performances. Then, she is equipped to give an account of the dimensions along which we can appropriately evaluate beliefs by referencing this performative aim. Moreover, on such an approach, the normative force of epistemic evaluation is relatively easy to explain: epistemic evaluation has normative force because beliefs constitutively aim at a certain goal. Perhaps for precisely this reason, many contemporary epistemologists seem to endorse this methodological approach to epistemic inquiry, at least in its broadest formulation.¹²

Additionally, the methodological approach I have developed here is naturalistically acceptable. Implementing this approach involves characterizing the constitutive aim of belief in a way that is directly responsive to empirical data regarding the ways in which the cognitive system operating within an organism serves to increase or maintain the biological fitness of the organism. As a result, this methodological approach sets empirical standards that any acceptable account of epistemic evaluation

¹² It seems to me, for example, that Nishi Shah and David Velleman adopt, at least in broad outline, something like the methodological approach that I have been developing here.

will have to meet. Specifically, according to the proposed methodology, any account of epistemic evaluation must characterize the constitutive aim of belief in a way that is, at the very least, consistent with empirical data regarding how the performances of the cognitive system serve to increase or maintain the biological fitness of the organism in which that system operates.

Furthermore, the sort of empirically driven approach I have outlined above is well-justified by its success in other domains. In particular, the strategy that the epistemologist taking on this methodology adopts in her inquiry is, in broad outline, the same as the strategy that the biologist or medical researcher adopts, with marked success, in her efforts to understand other organ systems in the human body (e.g. the circulatory system). Thus, the approach to epistemological inquiry I have outlined above serves to unify the project of developing an account of specifically epistemic evaluation—an account of the evaluation of the performances of our cognitive systems—with what I take to be similar projects in biology and medicine.

2. Implementing the Proposed Methodology: An Account of Epistemic Evaluation

In this section, I aim to articulate the account of epistemic evaluation that I argue results from implementing the methodology outlined above. Specifically, I propose here that in carrying out this methodological approach, one is led to conclude that the constitutive aim or goal of our cognitive performances—i.e. beliefs—is to be well-suited to serve as input to practical reasoning that yields successful action, and so beliefs are appropriately evaluated with reference to this aim. Before I defend my claim that implementing the methodological approach outlined above yields that the aim of our cognitive performances is to supply input to practical reasoning that leads to successful action, it will be helpful to spell out this proposal in greater detail.

First, let me try to make clear what it takes, on my view, for a belief to serve as input to practical reasoning. Practical reason, as I am conceiving of it here, is just the process through which beliefs interact with desires or preferences to generate intentions for action and so, generally, to generate actions themselves. Thus, for a subject's belief to serve as input to practical reason is for that belief to interact with desires or preferences such that the belief plays a role in shaping or exerts influence over the subject's subsequent intention(s) to act. A subject's belief that P will be well-suited to serve as input to practical reasoning to the degree that, if the belief that P were to play a role in shaping the subject's intentions and so her actions, the subject would be led to perform

actions which would be successful in normal worlds.¹³ The degree to which an action is successful, on my view, is just the degree to which it satisfies whichever desires or preferences the subject's intention to act was aimed to satisfy.¹⁴ Therefore, the degree to which a subject's belief is well-suited to serve as input to practical reasoning that yields successful action—and, thus, the degree to which a belief, understood as a cognitive performance, meets its aim—is just the degree to which that belief can inform action in a way that leads to the satisfaction of the subject's desires or preferences in a normal world.

2.1

Recall that, on the methodological approach to epistemological inquiry that I am adopting here, the first task facing the epistemologist developing an account of epistemic evaluation is the task of identifying the aim of our cognitive performances. Moreover, the epistemologist identifies the aim of our cognitive performances by first discovering how the performances of the cognitive system within an organism can and do contribute to the biological fitness of the organism. One discovers how any particular system operating within an organism contributes to the biological fitness of that organism by observing the ways in which the operation of the system in question enables the organism to cope with its natural environment. Thus, to determine the function of the cognitive system, and so the aim of cognitive performances, one must determine out how cognitive

¹³ Relativizing to normal worlds as I have done here is a standard externalist move (see, for example, Goldman, A. (1986), chapter 5). Generally, this move serves to free the reliabilist from being committed to the claim that brains in vats, for example, are never justified since none of their cognitive processes reliably generate true beliefs. Relativizing to normal worlds operates differently in the account of epistemic evaluation that I develop here, but the motivation for the move is similar.

¹⁴ This result follows from understanding practical reasoning as an organ system with the proper function of satisfying desires.

systems operating within organisms help these organisms cope with their natural environment in a way that enhances or maintains biological fitness.¹⁵

In implementing this methodological approach, one determines what it takes for the human being to negotiate her environment effectively (in a way that conducive to biological fitness) in order to discover how the cognitive system contributes to the human being's ability to cope with her environment. For an organism to achieve biological fitness in her natural environment, the organism's behavior must be calibrated to the various features of the environment. In cases where an organism's environment is uniform or predictable it is possible, and often quite effective, for the organism to be hard-wired (so to speak) to respond to particular stimuli with a fixed set of behaviors. For example, many animals grow winter coats as the days grow shorter since, in their natural environments, these animals need protection from the cold during the period of the year when the days are short. Having a winter coat is conducive to the biological fitness of these animals in their natural environments precisely because the growth of the winter coat is calibrated to various features of the environment (i.e. the number of hours of daylight). However, if the environment in which these animals were living were to undergo significant change of a certain sort—say the average temperature during the winter months rose —these animals would not be capable of adapting the growth of their winter coats to maintain biological fitness.

¹⁵ It is, of course, an empirical question how the operation of the cognitive system helps an organism cope with her environment, and so further empirical investigation may show that the answer to this question that I advance here falls short in certain respects. That said, I do believe and will make some effort to show that the empirical data here better supports the account that I develop in this section than the traditionally accepted view that cashes out the aim of the cognitive system in terms of truth.

The human environment, more so than other organisms' natural environment, is exceedingly variable and unpredictable. In order to survive and flourish in a variable and unpredictable environment, an organism must be able to select which behaviors to perform at any given time in a way that is *actively* and *flexibly* sensitive to or informed by the current state of her environment. So, an organism that achieves biological fitness in a variable and unpredictable environment must be equipped with a mechanism for selecting an appropriate behavioral response in a variety of different situations. Such a mechanism will have to take as input information about the environment and information about the organism's needs, and on the basis of this information generate a behavioral response to the situation at hand that helps the organism meet these needs given the current state of the environment. Practical reasoning, as I understand it, is just this mechanism. The cognitive system has the function or job of generating the first sort of input to practical reasoning and the conative system has the function of generating the second sort of input to practical reasoning.¹⁶

As I have characterized the mechanism of practical reasoning here, the aim of practical reasoning is to select a behavioral response to any given stimulus that will satisfy certain of the organism's desires. The cognitive system serves to supply representations of the current state of the organism's environment as input to the mechanism of practical reasoning. As a result, the proper function of cognitive system is to supply input to this mechanism that is well-suited to inform behavior in a way that,

¹⁶ Human desires do not always express biological needs the satisfaction of which will maintain or increase fitness. However, perhaps we should say that the function of the conative system, and so the aim of performances of this system (e.g. desires), is to represent or express needs of the organism, the satisfaction of which would, in fact, maintain or increase biological fitness.

given the current environment, leads the organism to satisfy her desires or needs. In turn, the performances of the cognitive system aim to serve as input to practical reasoning that yields successful action.

2.2

I have tried to show that implementing the methodological approach outlined above yields that the aim of our cognitive performances is to supply input to practical reasoning that engenders successful action rather than to represent truths. If I am right here, then, since, on the proposed methodology, performances are evaluated with respect to their aims, epistemic evaluation must be cashed out in terms of beliefs being well-suited to serve as input to practical reasoning that engenders successful action rather than in terms of truth. In what follows, I implement the second step of the methodological approach sketched above in an effort to outline such an account of epistemic evaluation.

Recall that, on the approach I am adopting here, all performances have AAA normative structure. A performance is accurate to the degree that the performance meets its aim. So, on the account of epistemic evaluation that I am developing here, a cognitive performance—a belief—is accurate to the degree that the performance is well-suited to serve as input to practical reasoning that yields successful action. A performance is adroit to the degree that it manifests skill. So, a cognitive performance will be adroit to the degree that it manifests cognitive skill. What it takes for a cognitive process or habit to count as a cognitive skill is for the cognitive process in question to reliably produce performances that meet their constitutive aim. Thus, on the view I am developing here, a cognitive process counts as a cognitive skill and so will be sanctioned by epistemic norms just in case it reliably produces beliefs that serve as input to practical reasoning that

yields successful action.¹⁷ So, a belief is adroit when it results from the operation of a cognitive process that reliably produces beliefs that serve as input to practical reasoning yielding successful action. Following Sosa, I suggest that justified belief is just adroit belief and so, on my view, a belief is justified when it is the product of a cognitive process that reliably produces beliefs that are well-suited to serve as input to practical reasoning that engenders successful action.¹⁸

Recall that a performance is apt when the performance is accurate because or as a result of being adroit. Thus, a cognitive performance is apt just in case it is well-suited to serve as input to practical reasoning that yields successful action as a result of being the product of a cognitive habit that reliably produces beliefs that are well-suited to serve as input to practical reason that issues in successful action. Again following Sosa, I propose that an account of apt belief serves as an account of knowledge.¹⁹ So, on my view, a belief counts as knowledge just in the case that belief is well-suited to serve as input to practical reasoning issuing in successful action as a result of having been produced by a cognitive process that reliably produces beliefs that meet this aim.

¹⁷ Reliability is, of course, a matter of degree here, although I will, for the sake of clarity and simplicity, generally gloss over this point in discussion that follows.

¹⁸ I am inclined to think that a belief is rational if and only if it is justified and so I would suggest that a belief is held rationally just in case the belief is adroit.

¹⁹ I suggest here that the human knowledge is properly characterized as apt belief (what Sosa identifies as animal knowledge). However, after identifying apt belief as animal knowledge, Sosa goes on to develop an account of a second and more demanding kind of knowledge: reflective knowledge. For Sosa, human knowledge is always reflective knowledge or apt belief aptly noted, and so is more demanding than mere animal knowledge or apt belief. I do not find the motivation for this move compelling. Thus, I resist Sosa's claim that human knowledge should be characterized as reflective knowledge or apt belief aptly noted.

2.3

It is worth pointing out that the account of the aim of our cognitive system I have developed here is more in line with many of our considered intuitions about cognitive performance than accounts cashing out the constitutive aim of our cognitive system in terms of truth. It is entirely consistent with my account that sometimes—perhaps even in the vast majority of cases—the cognitive system will have to generate beliefs that are true, approximately true, or true in the right respects in order to supply input to practical reasoning that yields successful action.²⁰ Nevertheless, it is a mistake to think that the aim of our cognitive performances is truth (or even approximate truth, or truth in the right respects). One reason this move is a mistake is that there may well be cases in which, perhaps because of the nature of our capacity for practical reasoning or perhaps because of certain features of our current environment, a belief that is false is particularly well-suited to serve as input to practical reasoning yielding successful action.

For example, Newtonian mechanics is, strictly speaking, false. Nevertheless, beliefs about how objects are likely to move that conform to Newtonian mechanics are well-suited to inform our actions in a way that engenders success. Perhaps the talented baseball outfielder is able to catch a fly ball because beliefs conforming to Newtonian mechanics inform his behavior. Similarly, the engineer employs beliefs conforming to Newtonian mechanics in designing a bridge that will remain structurally sound through an earthquake. Moreover, in each of these cases it seems plausible that, had the subjects' beliefs conformed instead with the (presumably true) general theory of relativity, the

²⁰ In fact, I suspect that it is rarely the case that beliefs that are completely true are well-suited to serve as input to practical reasoning. However, it may be that beliefs that are approximately true or true in the right respects are quite often extremely well-suited to serve as input to practical reasoning.

subjects would, perhaps because of the complexity of this theory, have been ill-equipped to act in a way that satisfied his or her preferences or desires. In each of these cases, it is intuitive to say that the beliefs in question, although false, are the *best* beliefs that the subject could have had. The account I am advancing of the aim of our cognitive performances explains *why* these beliefs are, in fact, the best beliefs that our subjects could have had. These beliefs, in each case, are incredibly well-suited to serve as input to practical reasoning that yields successful action. Now the epistemologist cashing out the aim of cognitive performances in terms of truth may well be able to explain our intuitions here. After all, Newtonian mechanics is approximately true or, perhaps, true in the right respects. However, the explanation that this epistemologist can offer, if successful at all, will certainly involve far more fancy philosophical footwork than the explanation that the account I am advancing here supplies.

Consider now the case of an athlete—a swimmer, let's say—who believes that she has a reasonable chance of winning a particular race. We can imagine, however, that the relevant empirical data strongly suggests that our athlete is unlikely to win the race (the other competitors in the race have consistently made better times in this event, the other competitors have been training just as hard for the race, etc.). Moreover, we can imagine that the swimmer maintains her belief that she has a reasonable chance of winning the race despite her being fully aware of the evidence suggesting otherwise. Research has shown that people who have somewhat unrealistic expectations regarding their own performance are overwhelmingly likely to perform better than matched peers with a more realistic assessment of their own potential. So, we can easily imagine that our swimmer, having convinced herself that she has a real shot at winning the race, goes

on to swim faster than she would have if she had believed, as the available evidence had suggested, that she was thoroughly unlikely to win.²¹ Perhaps, she wins the race after all or perhaps she merely swims a personal best time for that event or places a great deal higher than the relevant statistics suggested that she would. In any case, the athlete's false belief that she has a good chance of winning the race engenders successful action more effectively than the corresponding true belief would.

Moreover, the swimmer's false belief is, I suspect, precisely the belief that we think she *should* have regarding her own chances of winning. We would not, for example, encourage the swimmer to adopt a more realistic belief regarding her own chances of winning the race. Rather, we encourage precisely the sort of positive thinking that might have lead the swimmer to form her unrealistic belief in the first place. Not only do we think it is preferable, all things considered, that the swimmer have a somewhat unrealistic belief regarding her own chances of winning, but we are disinclined to say, in a manner that carries normative weight, that there is any respect in which the swimmer goes wrong in believing as she does.²²

I do not mean to suggest here that truth is always irrelevant to the epistemic status of a belief. In many cases, as a matter of contingent fact, it might well turn out that the sorts of beliefs that are best suited to serve as input to practical reasoning toward successful action are true beliefs. I merely suggest here that 1) it is a mistake to

²¹ Of course, if our swimmer's assessment of her own chances of winning are too unrealistic, then the belief is likely to be ill-suited to serving as input to practical reasoning yielding successful action since, for example, in such a case the swimmer will likely be incredibly disappointed by the results of the race.

²² Some epistemologists might be inclined to say that this swimmer is epistemically irrational. I hope that what I say below in Section 3.2 will make clear why it is a mistake to analyze the swimmer's case in this way.

characterize the aim of cognitive performances in terms of truth, and, as a result, 2) that there are, at least in principle, cases where there is nothing bad, epistemically speaking, to be said about an organism's holding a belief that is false. Put another way, on the view I am advancing here, it is, at least in principle, possible that there are false beliefs that, from an epistemic point of view, we *should*, nevertheless, hold, even in cases where a fair and thorough survey of the available evidence suggests to us that these beliefs are false. I take it that the swimmer's unrealistic estimation of her own chances of winning the race, for example, is a belief of just this sort. An account characterizing the aim of our cognitive performances in terms of truth is ill-equipped to make sense of these sorts of cases. Thus, one has good reason to reject the move to characterize the aim of belief in terms of truth and, instead, adopt a pragmatically-driven characterization of the aim of belief like the one I develop above.

More fundamentally, the move to characterize the aim of belief in terms of truth is a mistake because it signifies a failure to appropriately ground one's account of the aim of our cognitive performances, and so the normative force of epistemic evaluation in the particular job or function that the cognitive system operating within an organism performs in contributing to the biological fitness of that organism. I take the case of the swimmer with an unrealistic opinion of her own chances of winning to show that forming true, almost true, true-in-the-relevant-respects beliefs on a particular topic simply does not ensure that one's beliefs are well-suited to serve as input to practical reasoning toward successful action. If this claim is right, then adopting a characterization of the aim of belief in terms of truth is unmotivated by empirical considerations relevant to the proper function or purpose of the cognitive system and so the constitutive aim of our

cognitive performances. Thus, for the epistemologist who is committed to the methodology that I have adopted here, such a move involves an objectionable departure from the accepted methodological approach. If one accepts the methodology outlined above, one should accept the pragmatically-driven characterization of the aim of belief and the resulting account of epistemic evaluation that I have developed here over rival accounts that spell out the constitutive aim of belief in terms of truth.

2.4

Before moving forward, I wish to briefly highlight a feature of the account that I have developed here that I think counts in its favor: the account of epistemic evaluation that I have offered is contextualist. However, this account is contextualist in a way that differs from most traditional contextualist accounts of epistemic evaluation. Specifically, the aim of our cognitive performances, as I have characterized it here, is context-sensitive. Our cognitive performances aim to supply input to practical reasoning that yields successful action. However, I have cashed out successful action here in terms of desire or preference satisfaction and so the success of an action is, at least in part, determined by the particular preferences and desires of the organism performing the action. Of course, different organisms may well have different desires or preferences and a single organism's desires or preferences may change over time. So, whether or not a particular action counts as successful depends on facts about the actor at the particular time that she is performing the action. Thus, a fully specified characterization of the aim of the performances of any particular cognitive system will depend on facts about the preferences and desires of the organism within which that cognitive system operates at a certain point in time.

Most contextualist accounts build context sensitivity into their account of knowledge by building context sensitivity into their account of justification. Furthermore, traditional accounts often justify the context-sensitivity of their views simply by citing the fact that their contextualist accounts better capture how we speak and think about justification and knowledge in our everyday lives. The view on offer here builds context-sensitivity into the account of epistemic evaluation at a more fundamental level (at the level of characterizing the aim of our cognitive performances). As a result, both the account of accurate belief and the account of adroit belief (justified belief) I have developed here inherit a kind of context-sensitivity from the context-sensitivity built in to aim of our cognitive performances. An accurate belief is just a cognitive performance that meets its aim, and so whether or not a particular belief counts as accurate depends on the particular context-sensitive aim of that cognitive performance. Similarly, adroit belief is just belief that manifests cognitive skill. However, since a cognitive habit or process counts as a cognitive skill to the degree that it reliably produces cognitive performances that meet their aim, and since the aim of any particular cognitive performance is context-sensitive, whether or not a particular cognitive habit or process counts as a cognitive skill will also be somewhat context sensitive.²³ As a result, the view I advance shares the advantage of other contextualist accounts in vindicating our everyday way of thinking and speaking about knowledge and justification. However, the account of epistemic evaluation on offer here supplies a more principled justification for

²³ That said, I am sympathetic to the idea that cognitive skills and so epistemic norms are appropriately characterized in broad enough terms and at a high enough level of abstraction so that the context-sensitivity of the particular aim of our cognitive performances does not generate pervasive context-sensitivity in the epistemologist's account of epistemic norms.

its contextualism by locating context-sensitivity in the empirically discoverable aim of our cognitive performances.

3. Considering Objections

I have tried to show that if one adopts the methodology on offer here, then one should also adopt the account of epistemic evaluation that I have developed above. In what follows, I aim to defuse two related lines of objection to the account of epistemic evaluation I have sketched above. Both lines of resistance are, at their cores, attempts to argue that truth-based accounts of epistemic evaluation have some decisive, scale-tipping virtue that the pragmatically-driven account on offer here lacks. It is worth making explicit that I will not consider objections to the methodological approach that I have adopted here—I take commitment to this methodology as my starting point. Thus, my goal in defusing each of these objections is to bolster my claim that if one adopts the methodological approach outlined in Section 1, then one should prefer the pragmatically-driven account developed in Section 2 to an account of epistemic evaluation cashing out the aim of belief in terms of truth.

3.1

In “The Rationality of Belief and Some Other Propositional Attitudes”, Kelly defends the claim that practical considerations cannot rationalize belief.²⁴ Put in other terms, Kelly’s thesis is that the consequences of holding any particular belief do not and

²⁴ Kelly does not take himself here to have supplied a thorough defense of this claim, but rather to have laid out compelling support for this thesis and, by so doing, to have shifted the burden of proof to his opposition. In what follows, I hope to discharge this burden by describing how practical considerations play a role in rationalizing beliefs.

cannot directly affect whether it is rational for a subject to hold that belief.²⁵ At least at first blush, Kelly's thesis seems at odds with the account of epistemic evaluation that I developed above. This account grounds the normative force of epistemic evaluation in the constitutive aim of belief: beliefs aim to be well-suited to serve as input to practical reasoning toward successful action. Thus, according to the account on offer here, the source of the normative force of epistemic evaluations is characterized in thoroughly pragmatic or practical terms. As a result, since an account of rational belief is simply an account of a particular kind of epistemic evaluation, practical considerations will necessarily, on this account, figure prominently in an explanation of when and why a subject's belief that P counts as rational. In what follows, I aim to undermine Kelly's claim that practical considerations are irrelevant the rationality of any particular belief, and expose precisely where Kelly's defense of this claim fails, by articulating the particular role that practical considerations play in rationalizing beliefs in the account of epistemic evaluation developed above.

Kelly claims that only the grounds on which a belief is based can play a role in determining the rationality of that belief. Thus, for Kelly, whether or not a subject is rational in believing that P is settled by consideration of the particular grounds on which the subject's belief that P is based. If the subject believes that P on the basis of adequate grounds, then the subject is rational in believing that P. It is, on Kelly's view, precisely and exclusively the adequate grounds on which the subject's belief is based that rationalize her belief that P.

²⁵ Of course, Kelly admits that the expected consequences of holding a particular belief may figure in a complete story detailing the causal history of the belief.

Kelly goes on to argue that beliefs cannot be based on practical considerations. In so doing, Kelly suggests that one can determine the considerations upon which a subject bases her belief that P by discovering what sorts of considerations might prompt the subject to give up her belief that P. In Kelly's words, "[t]he considerations on which a given belief (or course of action) is based are revealed by the circumstances which would prompt one to abandon that belief (or course of action)."²⁶ Kelly goes on to point out that, at least as a matter of empirical fact, becoming aware of certain likely consequences of believing that P does not prompt a subject to revise her credence in P or abandon her belief that P. Becoming aware of certain likely consequences of believing that P may prompt the subject to take action aimed at indirectly influencing her belief that P. However, a subject will not revise or abandon her belief that P directly in response to a change in the relevant practical considerations. Kelly takes this fact as strong evidence for his claim that beliefs cannot be based on practical considerations. Thus, since on Kelly's view, the nature of the grounds on which a belief is based settle the questions of whether the belief is rationally held, Kelly concludes that practical considerations cannot play a role in rationalizing belief.

In brief outline, Kelly's argument for the claim that practical considerations cannot rationalize belief is as follows:

Premise 1: The nature of the grounds on which a subject's belief that P is based settles the question of whether the subject rationally believes that P.

Premise 2: Practical considerations cannot serve as the grounds on which a subject's belief that P is based.

Conclusion: Practical considerations play no role in rationalizing belief.

²⁶ Kelly, T. (2002), 176.

If Kelly is right that practical considerations cannot serve as the grounds on which a subject bases her belief and that the nature of a subject's grounds for belief settles the question of whether to subject's belief is rational, then follows that practical considerations cannot directly affect whether it is rational for a subject to believe that P. However, in what follows I hope to show that Kelly is wrong to think that only the grounds on which a subject bases her belief can rationalize the belief.²⁷ In so doing, I aim to undermine Kelly's conclusion that practical considerations do not impact the rationality of any particular belief. Thus, although Kelly may well be right in claiming that practical considerations cannot serve as the grounds on which a subject bases her belief, I aim to show that practical considerations do, nevertheless, play a crucial role in rationalizing beliefs.

Recall that, on the view I have been advancing here, a belief is rationally held if and only if the belief is adroit or manifests a kind of cognitive skill. That is, a subject rationally believes that P just in case her belief is the product or output of a cognitive process that reliably generates beliefs that meet their performative aim. So, a subject's belief that P is rational just when the belief is the product of a cognitive process that reliably produces beliefs that are well-suited to serve as input to practical reasoning issuing in successful action. In order to determine whether or not a particular belief is rational, then, one must first discover which cognitive process gave rise to the belief, and then check to see whether or not the cognitive process in question is one that reliably generates beliefs well-suited to serve as input to practical reasoning yielding successful

²⁷ I suspect that one might well resist the particular account of epistemic evaluation I have advanced above and yet, nevertheless, successfully adopt precisely this strategy in objecting to Kelly's argument here.

action. A cognitive process is individuated, at least in part, by which inputs or types of input the process takes as arguments—sorts of grounds upon which the output (belief) will be based. So, identifying the particular cognitive process that generates a given belief involves identifying the inputs or grounds upon which that belief is based. As a result, the grounds on which a subject's belief that P is based play a crucial role in rationalizing the subject's belief. However, facts establishing that the cognitive process generating the subject's belief that P is a reliable one also play an equally crucial role in rationalizing the subject's particular belief that P. Determining whether the cognitive process that gives rise to a particular belief reliably generates beliefs that meet their performative aim involves checking to see if beliefs formed on the basis of a certain sort of input are likely to guide action successfully. That is, the consequences of forming beliefs in a particular manner—the consequences of employing a particular cognitive process to form beliefs in a certain domain—determine whether beliefs formed accordingly are well-suited to serve as input to practical reasoning guiding successful action. Thus, the consequences of forming beliefs through the operation of a particular cognitive process dictate whether forming beliefs in the given domain through the operation of this cognitive process manifests a kind of cognitive skill. Paradigmatically practical considerations—information regarding the consequences of forming beliefs in a given domain through the operation of a certain cognitive process—play a crucial role in determining whether or not a particular belief that P, formed through the operation of the cognitive process in question and so on the basis of certain sorts of grounds, is rationally held. More succinctly, practical considerations play a role in rationalizing belief by determining when the grounds on which a belief is based count as adequate. Therefore,

on the view of epistemic evaluation I have been advancing here, practical considerations do figure prominently in an account of when and why a particular belief is rational even if, as Kelly suggests, they cannot ground beliefs.

3.2

There is a second intimately related and perhaps more fundamental line of resistance to the account of epistemic evaluation developed above that I wish to address here. The objector adopting this line of resistance maintains that any account of epistemic evaluation must be capable of sustaining a substantive distinction between two different normative domains: the epistemic domain and the practical or pragmatic domain. She goes on to suggest that the account of epistemic evaluation I am advancing here fails to carve out an independent, distinctively epistemic domain of normative evaluation and, instead, merely subsumes epistemic evaluation as a species of pragmatic or practical evaluation. Put another way, the objector claims that the account of epistemic evaluation that I am advancing here lacks the resources to underwrite or vindicate a substantive distinction between epistemic evaluation and practical evaluation.

In fact, I suspect it is precisely this line of thinking that motivates Kelly's efforts to defend the claim that practical considerations cannot rationalize belief. In his introduction, Kelly makes explicit that he is concerned "with the fact that one can evaluate one's own beliefs practically as well as epistemically."²⁸ Additionally, I suspect that it is precisely this line of thought that motivates certain epistemologists to characterize the performative aim of beliefs in terms of truth rather than in pragmatic

²⁸ Kelly (2002), 164.

terms as I have proposed.²⁹ It seems that epistemologists of this mind count the fact that accounts cashing out the aim of belief in terms of truth straightforwardly carve out an independent domain of distinctively epistemic evaluation as a decisive virtue of such accounts. My aim here is to undermine this line of resistance to the pragmatically-driven account of epistemic evaluation I have been advancing here by showing that this account does, in fact, have the resources to mark out a domain of distinctively epistemic evaluation.

The source of this objector's worry that the account of epistemic evaluation on offer here falls short in this respect is that, according to this account, the source of the normative force of epistemic evaluations is, at its core, pragmatic in character. Recall that on the methodological approach I have adopted here, it is the constitutive aim of belief that gives epistemic evaluation—evaluation with respect to this aim—normative force. Moreover, as I have characterized it above, the constitutive aim of belief is a paradigmatically pragmatic one. Thus, the account of epistemic evaluation I am advancing here explains the normative force of our epistemic evaluations by pointing to the fact that performances of the cognitive system necessarily aim to be well-suited to serve as input to practical reasoning generating successful action. At least at first blush, then, the objector's accusation seems to stick. It seems that, on the account I have been advancing here, epistemic evaluation is, at bottom, merely a particular species of practical evaluation.

²⁹ I am thinking here, in particular, of Shah and Velleman, and, although it is more difficult to tell, perhaps Sosa as well.

Perhaps it will be helpful in fleshing out the objector's line of resistance to revisit the case of the overly optimistic athlete.³⁰ Recall again the swimmer who, in the face of evidence suggesting otherwise, believes that she has a reasonable shot at winning a race. Even if one accepts that it is better, all things considered, for the swimmer to believe as she does, the objector claims that one can and should still maintain that the swimmer's belief is epistemically unjustified and/or epistemically irrational. In developing the account of epistemic evaluation that I have been advancing here, of course, I denied precisely this claim. I suggested that there is no sense in which the swimmer does something wrong in believing as she does. That is, the swimmer's overly optimistic belief, I maintained, is not a belief the holding of which under the specified circumstances merits *any* sort of criticism leveled with genuine normative force. Thus, the account I have been advancing here cannot support the claim that the swimmer's belief is epistemically unjustified, although, perhaps, practically rational and, all things considered, the belief that she should hold under the circumstances. Put another way, this pragmatically-driven account cannot vindicate an analysis of the swimmer's case maintaining that the swimmer's belief is epistemically faulty, but practically sound. It is this sort of observation, I suspect, that prompts the objector to resist the account of epistemic evaluation I have been advancing here on the grounds that this account lacks the resources to underwrite a substantive distinction between epistemic and practical evaluation.

³⁰ Kelly, in fact, briefly discusses precisely this case in setting up his own view (see Kelly (2002), 164). Kelly suggests here that the athlete's overly optimistic belief is irrational, although perhaps practically sound.

I will argue that analyzing the swimmer's overly optimistic belief as epistemically faulty, but practically sound, is deeply problematic. After all, I have suggested that it is an advantage of the account that I have been advancing here that it avoids commitment to this sort of analysis of the swimmer's overly optimistic belief. Furthermore, I hope to show that the objector is overly hasty in concluding, on the grounds that the account of epistemic evaluation on offer here cannot vindicate such an analysis of the swimmer's overly optimistic belief, for example, that this account lacks the resources to mark out an independent and distinctively epistemic domain of normative evaluation.

Before articulating the way in which the account I have been advancing here circumscribes a normative domain of distinctively epistemic evaluation, I will briefly suggest why I find the sort of analysis of the swimmer's overly optimistic belief maintaining that the swimmer's belief is epistemically faulty unpalatable. Recall the intuitive plausibility of the claims that in holding her overly optimistic belief, the swimmer believes exactly as she should and that the swimmer does nothing wrong and so does not deserve criticism or reproach in believing as she does. These claims have normative force. That is, this way of evaluating the swimmer's overly optimistic belief guides the way we think and act in responding to the case. Not only do we refrain from criticizing the swimmer for believing as she does, but we often actively encourage her to maintain her overly optimistic belief in the face of defeating evidence. So, if, as the objector here suggests, the swimmer's overly optimistic belief is epistemically faulty and merits criticism, then the epistemic status of the swimmer's belief carries little or no normative weight. Analyzing the swimmer's overly optimistic belief as epistemically faulty, but pragmatically sound, robs the epistemic evaluation of the swimmer's case of

any noteworthy normative or action-guiding force. So, if the objector is right in claiming that the swimmer's belief is epistemically faulty, then it follows that our epistemic evaluation of the swimmer's belief is impotent with respect to informing or influencing how we act in response to the swimmer's case. After all, we are not moved to act as if there is something wrong with the swimmer's believing as she does. Thus, I suggest that one should reject analyses of the swimmer's overly optimistic belief maintaining that the belief is epistemically faulty on the grounds that such analyses undercut the normative force of epistemic evaluation.

It is a fact that the pragmatically-driven account of epistemic evaluation I have been defending here will not vindicate the claims like the claim that the swimmer's overly optimistic belief is epistemically faulty. However, contra the objector adopting the line of resistance described above, I maintain that this fact is not symptomatic of the account's inability to vindicate the platitudinous thought that the epistemic evaluation is distinct from a kind of practical evaluation. In what follows, I sketch an explanation that an epistemologist who wishes to take on what I have proposed above might flesh out in an effort to show that her pragmatically-driven account of epistemic evaluation can underwrite a substantive distinction between epistemic and practical evaluation. My aim here is not to develop this explanation in detail, but rather to shift the burden of proof to the objector who endeavors to push this line of resistance against this account.

Recall that, on the account of epistemic evaluation I have been advancing here, to engage in epistemic evaluation is just to evaluate the performances of the cognitive system with respect to their particular constitutive aim. So, at the very least, the objects of epistemic evaluation are well-defined and distinctive: epistemic evaluation evaluates

the performances of the cognitive system (i.e. beliefs). Moreover, epistemic evaluation is necessarily evaluation with respect to the particular performative aim of these performances. That is, epistemic evaluation is necessarily evaluation with respect to whether a performance is (accuracy) or is likely to be (adroitness) well-suited to serve as input to practical reasoning yielding action that fulfills certain desires or needs. Thus, as I have argued above, pragmatic considerations will necessarily play some role in an account explaining when and why certain of our beliefs have positive epistemic status.³¹

However, there is no reason to think that pragmatic considerations will play an analogous role in an account of practical evaluation. The epistemologist adopting the account of epistemic evaluation that I have been advancing here can characterize practical evaluation as, at least in the first instance, evaluation of actions (or, perhaps, intentions to act). On the account that I have been advancing here, actions are appropriately understood as the output of practical reasoning and so an action is successful or meets its aim precisely when it satisfies those desires or needs that motivate the action (i.e. those desires or needs that served as input to the bit of practical reasoning generating the action in question as output). So, engaging in practical evaluation involves evaluating an action with respect to whether the action, in fact, satisfies those desires that motivated the performance of the action, or whether the action is likely to satisfy these desires. As a result, the norms governing practical evaluation will be crucially different in character from the norms governing epistemic evaluation. Norms governing epistemic evaluation—and, in particular, adroit belief—will, roughly speaking, have the following form: employ cognitive process *p* (specified, in part, by the sorts of

³¹ I hope that what I have said in 3.1 helps to make clear what this role will be.

grounds on which beliefs generated by this cognitive process are based) in circumstances c to form beliefs that are suitable to serve as input to practical reasoning with respect to domain z . However, norms governing practical evaluation—and, in particular, adroit action—will, again, roughly speaking, have the following alternate form: if one has beliefs $b_1...b_n$, then perform an action of type a to satisfy desires of type d . Thus, epistemic norms specify when a particular sort of performance (i.e. a belief) is likely to lead an organism to perform an action that satisfies certain of its desires, while practical norms specify when an organism's performance (in the first instance, an action) is likely to satisfy whichever of the organism's desires motivate that performance. On this account, both epistemic and practical evaluation is, at bottom, pragmatically-driven. However, practical considerations play crucially different roles in specifying epistemic and practical norms.

I have emphasized that practical evaluation, according to this view, is, at least in the first instance, evaluation of actions or intentions to act. Nevertheless, one can easily imagine using practical norms to evaluate beliefs.³² When one does evaluate a belief with respect to both practical and epistemic norms, it is entirely possible for the epistemic and the practical evaluations of the belief to come apart. Imagine that a soldier has gone missing and, after some time, the military contacts the soldier's wife to inform her that her husband is missing and presumed dead. Upon hearing this news, it seems reasonable to assume that either the soldier's wife will form the belief that her husband is dead or the soldier's wife will come to believe that her husband is missing and possibly still alive. On the account of epistemic evaluation that I have been advancing here, if the wife comes

³² Although, doing so would, perhaps, be misguided in certain respects.

to believe that her husband is dead, then her belief will have positive epistemic status. One can easily imagine that the wife's belief that her husband is dead is formed in a way that reliably generates beliefs that are well-suited to serve as input to practical reasoning issuing in successful action. After all, forming beliefs on the basis of expert testimony is plausibly quite reliable in this respect. Thus, the wife's belief that her husband is dead is adroit. Moreover, it is plausible that the wife's belief that her husband is dead is, itself, well-suited to serve as input to practical reasoning toward successful action and so the belief is accurate. So, from an epistemic perspective, one must say that the wife *should* believe that her husband is dead. However, it is plausible that, from the practical perspective I have sketched above, it would be best for the wife to believe that her husband is possibly still alive since holding this belief, one can easily imagine, would satisfy the wife's immediate desires far more effectively than would holding the belief that the soldier is dead. In particular, it is plausible that the wife wants to believe that her husband is alive and, in fact, it is easy to imagine that the wife comes to believe that her husband is still alive precisely because so believing satisfies this desire. That is, one can easily imagine that the wife's belief is the result of wishful thinking. If, as I have suggested, this is a case in which epistemic evaluation recommends believing one thing—namely, that the soldier is dead—and practical evaluation recommends believing otherwise—namely, believing that the soldier is possibly still alive—then it should be clear that the account in question can underwrite a substantive distinction between epistemic and practical evaluation.

I have tried to sketch a way in which an epistemologist adopting the account of epistemic evaluation that I have been advancing here might vindicate the intuition that

practical evaluation and epistemic evaluation are fundamentally different and so can, at least in some case, yield divergent recommendations with respect to what one should believe. I hope that what I have said here is sufficient to undermine the objector's claim that the account on offer here lacks the resources to circumscribe a normative domain of distinctively epistemic evaluation. If I have been successful in this effort, then the epistemologist who endorses the methodology outlined above should not resist adopting the pragmatically-driven account of epistemic evaluation that I have been advancing here instead of an account that cashes out the constitutive aim of belief in terms of truth.

Conclusion

In this paper, I have tried to show that if one adopts the methodological approach to epistemic inquiry outlined in Section 1, then one should also adopt the pragmatically-driven account of epistemic evaluation developed in Section 2 and defended in Section 3. The methodology on offer here understands beliefs as performances of the cognitive system and so maintains that account of epistemic evaluation—evaluation of beliefs—should be grounded in an account, given in general terms, of the way in which performances are appropriately evaluated. According to this methodological approach, performances are appropriately evaluated with reference to their constitutive aims. Thus, the epistemologist adopting this approach develops an account of epistemic evaluation by referencing the constitutive aim of belief. On the methodological approach on offer here, the constitutive aim of the performances of an organ system operating within an organism is derived from the empirically discoverable proper function of that system. I have argued here that empirical data supports the claim that the proper function of the cognitive system is to provide input to practical reasoning in a way that engenders successful action and so beliefs—understood as performances of the cognitive system—aim to be well-suited to serve as input to practical reasoning yielding successful action. I employed this characterization of the aim of belief in developing an account of epistemic evaluation here that cashes out accurate, adroit, and apt belief in pragmatic terms, and then tried to defuse two related lines of resistance to the resulting pragmatically-driven account of epistemic evaluation by articulating the particular role that practical

considerations will play in specifying epistemic norms and showing that the account on offer here has the resources to circumscribe a normative domain of distinctively epistemic evaluation. In so doing, I have hoped to establish that the pragmatically-driven account of epistemic evaluation developed above, rather than accounts that cash out epistemic evaluation in terms of a truth-goal, is the account that results if one rigorously and carefully put into practice the methodology on offer here. If I have been successful, then the epistemologist must either reject this methodological approach to epistemic inquiry or accept the account of epistemic evaluation that I have developed and defended here.

References

- Goldman, A. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Kelly, T. (2002). The rationality of belief and some other propositional attitudes. *Philosophical Studies*, 110, 163-96.
- Plantinga, A. (1993). *Warrant and Proper Function*. Oxford: Oxford University Press.
- Shah, N. & Velleman, J.D. (2005). Doxastic deliberation. *The Philosophical Review*, 114(4), 497-534.
- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford: Clarendon.
- Sosa, E. (2001). Reflective knowledge in the best circles. In M. Steup (Ed.) *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* (pp. 187-203). Oxford: Oxford University Press.