

Epistemic Rationality as Instrumental Rationality: A Critique

THOMAS KELLY

University of Notre Dame

1. Introduction

My aim in this paper is to explore the relationship between epistemic rationality and instrumental rationality. By *epistemic* rationality, I mean, roughly, the kind of rationality which one displays when one believes propositions that are strongly supported by one's evidence and refrains from believing propositions that are improbable given one's evidence. Prominent epistemologists frequently emphasize the disparate ways in which this term is employed and occasionally question its theoretical usefulness on this account.¹ With an eye towards such concerns, I will in what follows consider only examples in which the correctness of its application is more or less uncontroversial. Thus, if I have strong, undefeated evidence that the butler committed the crime, and my belief that the butler committed the crime is based on that evidence, then my belief that he did so is epistemically rational. By *instrumental* rationality, I mean the rationality which one displays in taking the means to one's ends. Thus, if I have the goal of asking the speaker a question, and I know that I will only be able to ask the speaker a question if I raise my hand, then (all else being equal) it is instrumentally rational for me to raise my hand.

How are epistemic and instrumental rationality related? Here is a particularly radical suggestion: epistemic rationality *just is* instrumental rationality. More precisely: epistemic rationality is a species of instrumental rationality, viz. instrumental rationality in the service of one's *cognitive* or *epistemic* goals. Call this way of thinking about epistemic rationality the **instrumentalist conception** of epistemic rationality. My primary concern in this paper is to explore the instrumentalist conception of epistemic rational-

¹ Plantinga (1993) distinguishes five 'varieties' of rationality; Goldman (1986) explicitly excludes rationality from the terms of epistemic evaluation which he seeks to analyze on the grounds that 'this notion is so vague in ordinary usage, and so disparately employed by different philosophers and social scientists, that it has limited usefulness' (p. 27).

ity—what is involved in thinking about epistemic rationality in this way, why this view would be of philosophical importance if true, and whether it is true or false. I will argue that although it possesses a certain intuitive appeal and enjoys considerable popularity among both epistemologists and philosophers of science, the instrumentalist conception is ultimately indefensible. After having argued for the distinctness of epistemic rationality and instrumental rationality, I will in a final section of the paper attempt to delineate the role of each in typical instances of theoretical reasoning.

First, some clarification. On anyone's view, the fact that I possess certain cognitive goals can make it instrumentally rational for me to do things which it would not be instrumentally rational for me to do, if I did not possess those goals. Suppose that, wanting to know the identity of the person who committed the crime, I engage in the activity of *looking for evidence* which bears on the question. Here, the fact that I have the goal of learning a certain truth gives me an instrumental reason to act in a certain way: all else being equal, it is rational for me to engage in the activity of looking for evidence. Uncontroversially, the rationality in play here is instrumental rationality in the service of a cognitive goal. Suppose that my search is successful: I discover strong evidence that the butler committed the crime. The character of this evidence singles out a certain response on my part as the epistemically rational response: it is rational for me to believe that the butler committed the crime. What is the relationship between the rationality which I exhibit in responding to the evidence in the epistemically appropriate way, and the rationality which I exhibit in acting so as to acquire that evidence? As we will see, this question *is* controversial. An instrumentalist wants to assimilate the rationality of my responding to the evidence in the epistemically appropriate way to the rationality of my looking for that evidence in the first place. Those who reject the instrumentalist conception, on the other hand, think that it is a fundamental mistake to think about epistemic rationality in this way.

That I have the goal of asking a question gives me a reason to raise my hand; that I have the goal of avoiding the flu gives me a reason to get a flu shot. But no one would think that there is some deep distinction between two kinds of rationality here: asking-a-question rationality and avoiding-the-flu rationality. On the other hand, some have thought that there is a deep and fundamental distinction between epistemic rationality and other types of rationality. If the instrumentalist conception of epistemic rationality is correct, however, then this thought is mistaken, and it is mistaken in exactly the same way as the thought that there is some fundamentally different kind of rationality called asking-a-question rationality. The instrumentalist conception is thus at bottom a reductionist view: it entails that there is, in fact, only one thing where it is natural to suppose that there are two.

The instrumentalist conception enjoys considerable popularity among both epistemologists and philosophers of science. It is, for example, the guiding idea behind one of the most thoroughly developed and theoretically sophisticated theories of epistemic rationality to be put forth in recent decades, Richard Foley's 'subjective foundationalism'.² For Foley, all rationality—the rationality of belief as well as the rationality of action—is a matter of rationally pursuing one's goals. According to Foley, epistemic rationality is distinguished from other types of rationality simply by its distinctive goal: the goal of now believing true propositions and not now believing false propositions.³

Within the philosophy of science, the instrumentalist conception is endorsed by Larry Laudan, author of a much-discussed position known as 'normative naturalism'. According to Laudan

Epistemic rationality...is simply a species of the genus instrumental rationality...Epistemic rationality, no less than any other sort of rationality, is a matter of integrating ends and means...Good reasons are instrumental reasons; there is no other sort (Laudan, 1990b, p. 318).

A list of other prominent philosophers who have explicitly expressed enthusiasm for this way of thinking about epistemic rationality would include Robert Nozick (1993, ch.3), Philip Kitcher (1992), and Ronald Giere (1989).

Before inquiring as to the correctness of the instrumentalist conception of epistemic rationality, I want to take up the question of *why it matters* whether this view is true or false. Why might someone want this view about epistemic rationality to be true?

2. The Instrumentalist Conception: Why It Matters

2.1 *The Instrumentalist Conception and Naturalism*

For Laudan, the instrumentalist conception of epistemic rationality is central to the project of naturalizing epistemology and the philosophy of science while preserving their normativity (Laudan 1996, ch.9). The essential idea is due to Quine. In his "Reply to Morton White", Quine wrote

Naturalization of epistemology does not jettison the normative...For me, normative epistemology is a branch of engineering. It is the technology of truth-seeking...it is a matter of

² As presented in his (1987) book.

³ Foley (1987, Ch.1. See especially pages 6-8.). Foley is also read in this way by both Plantinga (1993, p. 27) and Harman (1999b, p. 101). In later work (e.g., 1993), Foley sometimes characterizes the epistemic goal as that of 'having an accurate and comprehensive system of beliefs'. This difference is immaterial to the discussion which follows.

efficacy for an ulterior end, truth...The normative here, as elsewhere in engineering, becomes descriptive when the terminal parameter is expressed.⁴

In general, the idea that the normativity of epistemology is simply the normativity of instrumental reason is especially popular among those who, following Quine, advocate the naturalization of epistemology and the philosophy of science but who do not want to abandon the traditional normative aspect of those disciplines.⁵

It is not difficult to see why the instrumentalist conception of epistemic rationality would be popular from the perspective of naturalism. For it is widely held, by both enthusiasts for and detractors of naturalism in philosophy, that the apparent existence of various kinds of normativity constitutes one of the greatest potential obstacles for naturalism.⁶ The burden of the naturalist is thus to show that any apparent kind of normativity is either spurious or naturalistically unproblematic. With respect to epistemic normativity, several of the options available to the naturalist are fairly radical. Thus, a naturalist might be an *eliminativist* about epistemic normativity and advocate the replacement of normative epistemology by a purely descriptive branch of cognitive psychology. Quine is often read as such an eliminativist about epistemic normativity.⁷ Alternatively, a naturalist might offer a non-cognitivist, *expressivist* account of epistemic normativity, according to which claims about what it is epistemically rational to believe are neither true nor false, but merely serve to express the attitude of the speaker towards the norms which license the belief in question. Hartry Field—whose career has largely been devoted to the project of naturalizing that which seems beyond the naturalist pale—has recently embraced expressivism in epistemology.⁸

⁴ Quine (1986, pp. 664-665). Quine is in the course of explaining to Morton White why, contrary to what White and many others had supposed, Quine's persistent calls for a naturalized epistemology are *not* calls for doing away with normative epistemology. Compare Quine's remarks in his later (1993, p. 19).

⁵ In addition to Quine and Laudan, a list of philosophers who endorse this conception of epistemic normativity as a means to naturalizing epistemology would include Hilary Kornblith (1993), Kitcher (1992), Giere (1989), and James Maffie (1990a, 1990b). Foley's enthusiasm for this way of thinking about epistemic rationality does *not* seem to be rooted in naturalist concerns.

⁶ For a recent argument that the existence of normativity undermines naturalism, see Parfit (forthcoming).

⁷ For example, by Jaegwon Kim (1993). But this, as we have noted, is a misreading of Quine—although perhaps an understandable misreading, given some of Quine's early pronouncements. If in fact a thoroughgoing naturalist is ultimately committed to eliminativism about epistemic normativity, then this would be, as Frank Jackson has said, 'strong beer' (1999, p. 434).

⁸ In his (2000). Perhaps the first philosopher to explicitly consider expressivism in epistemology was Roderick Chisholm (1957). (But for Chisholm's views on normativity in epistemology, see also note 12 below.)

In contrast to such radical alternatives, the instrumentalist conception of epistemic rationality seems to promise a way of preserving a full-blooded, cognitivist account of epistemic normativity which is naturalistically unproblematic. After all, many philosophers regard the normativity characteristic of the reasons which one has to take the means to one's ends as utterly unproblematic for naturalism. And if in fact the normativity of instrumental reason is naturalistically unproblematic, and epistemic normativity is simply the normativity of instrumental reason, then (presumably) epistemic normativity is itself naturalistically unproblematic. The truth of the instrumentalist conception of epistemic rationality then, would seem to be something of a *coup* for the naturalist. Thus, it is not surprising that the assimilation of epistemic rationality to instrumental rationality should often be viewed with great enthusiasm by proponents of naturalism.⁹

Conversely, those who have attacked this conception of epistemic rationality have typically been staunch opponents of naturalism. Again, the motivation for such attacks is not hard to discern. In particular, the vindication of the instrumentalist conception would seem to undermine a favorite tactic of opponents of naturalism, viz. the appeal to 'companions in the guilt' arguments. It is widely thought, by both friends and foes of naturalism, that the existence of anything which possesses *categorical normative force*—that is, force which is binding on any rational agent, regardless of the goals or ends which he or she happens to hold—is not a possibility which the naturalist world view countenances.¹⁰ Of course, categorical normative force is exactly the kind of force which *moral* reasons are often claimed to possess. In view of this, many naturalists are quite prepared to dispense with moral reasons so construed—for such thinkers, the fact that naturalism does not countenance the existence of such reasons no more counts against naturalism than the fact that naturalism refuses to countenance, say, divine intervention in human affairs. It is at this point that the defender of moral reasons is apt to appeal to a 'companions in the guilt' argument and remind the naturalist that epistemic reasons, no less than moral reasons, seem to have categorical normative force. And because many naturalists who would not hesitate to throw out moral reasons would hesitate to throw out epistemic reasons, this is indeed a powerful rejoinder by the opponents of naturalism.¹¹

⁹ Although I have here presented the instrumentalist conception as an *alternative* to expressivism in epistemology, it's worth noting that expressivism in epistemology is in fact compatible with the instrumentalist conception: one might hold that epistemic rationality is instrumental rationality, and then proceed to tell an expressivist story about instrumental rationality. On the other hand, one might be an expressivist about epistemic rationality while rejecting the instrumentalist reduction.

¹⁰ A particularly clear and prominent statement of this thought is Mackie (1979). See especially Ch. 1, 'The Subjectivity of Values'.

¹¹ Hilary Putnam is among the most prominent of those who have attempted to tie the fate of moral reasons to epistemic reasons in an effort to defend the former. See, e.g., his

The instrumentalist conception of epistemic rationality seems to threaten this otherwise-powerful rejoinder by showing that epistemic reasons are *not* companions to moral reasons in the relevant respect: contrary to what one might have thought, epistemic reasons get their grip on us only insofar as we possess certain cognitive goals. The normative force of epistemic reasons is not, after all, categorical, but rather hypothetical. The triumph of the instrumentalist reduction would seem to show that one can throw out any alleged entities with categorical normative force *without* dispensing with epistemic reasons. Moral reasons might not have any companions in the guilt. There is then, a strong incentive for the opponent of naturalism to show that the instrumentalist conception of epistemic rationality is mistaken.¹²

To this point, much of what little explicit discussion the instrumentalist conception has received has taken place within the context of larger debates over the tenability of naturalism.¹³ However, the interest of the instrumentalist conception is not, I want to insist, exhausted by its potential implications for the project of naturalizing the normative. Suppose that it turns out that, contrary to what many assume, the normativity involved in taking the means to one's ends is *not* naturalistically unproblematic, and that, moreover, there is no naturalistically acceptable account of instrumental rationality to be had.¹⁴ If that turned out to be the case, then clearly, the envisaged reduction would hold little if any appeal for the naturalist. Nevertheless, the instrumentalist conception of epistemic rationality would still be an interesting view, for it is, I believe, a philosophically interesting view in its own right. If in fact epistemic rationality turns out to be a special case of instrumental ration-

(1990). As one would expect, Putnam is also a critic of naturalism in epistemology (Putnam, 1983). Compare Derek Parfit:

If moral reasons were too queer to be part of the fabric of the Universe, that would be true of all normative reasons, including reasons for believing. That conclusion is incredible...If moral skeptics wish to avoid such all-embracing skepticism, they must abandon these objections to moral realism. If reasons for believing are not incompatible with a scientific world view, nor are...[moral reasons] (forthcoming, p. 29).

The strategy of defending moral reasons by tying their fate to that of epistemic reasons has also been pursued by Frank Jackson (1999). For stimulating discussion of related issues, see also David Velleman (2000a) and Peter Railton (1997).

¹² Of course, even if it turns out that the normativity of epistemology is not reducible to the normativity of instrumental reason, this wouldn't show that epistemic normativity is irreducible, or that it is (as one says) 'sui generis'. One of the most prominent epistemologists of the twentieth century, Roderick Chisholm, was a longtime advocate of the interesting if eccentric view that epistemic normativity is really a species of *ethical* normativity. See, e.g., his (1991, p. 119) where he notes his career-long disagreement with Roderick Firth on this issue. For Firth's side of the argument, see his (1998a) and (1998b).

¹³ I have in mind here especially the exchanges between Laudan (1996), Siegel (1989, 1990, 1996), and Giere (1989).

¹⁴ For arguments that this is in fact the case, see Korsgaard (1997), Hampton (1998, especially Part 2, "Instrumental Reason") and Parfit (forthcoming).

ality, then this would be a deep and unobvious fact about the nature of epistemic rationality—and therefore, a fact of considerable interest for the epistemologist.

Consider an analogous case drawn from the philosophy of mathematics. Like instrumentalism, *logicism* is a reductionist thesis: roughly, logicism is the thesis that mathematical truth is really just logical truth. In the present century, much of the enthusiasm for logicism has been on the part of empiricists. It's not hard to see why logicism might look attractive to an empiricist: given that mathematics has always been the great thorn in the side of empiricism, the suggestion that mathematical truth is reducible to some other kind of truth looks like progress, or at least, potential progress. Of course, even if the logicist reduction had gone through, it's not as though the empiricist would have been home free. In particular, the empiricist would still have been faced with the task of showing why empiricism is not undercut by logic—surely no easy task. Now, it might be that there is no satisfying empiricist story to tell about logic; and in that case, the distinctly empiricist motivation for the logicist program would be undercut. Even so, it would be a great mistake to conclude that *logicism* is therefore devoid of interest. On the contrary, if mathematical truth had turned out to be reducible to logical truth, then this would be an extremely interesting fact about the nature of mathematical truth, even if a fact which is irrelevant to the traditional debate between empiricism and rationalism. Analogously, if epistemic rationality is reducible to instrumental rationality, then this would be an extremely interesting fact about the nature of epistemic rationality, even if a fact which is irrelevant to the ongoing debate over the merits of naturalism.¹⁵

Moreover, in addition to its intrinsic interest, whether the instrumentalist conception ultimately proves tenable may very well have important implications for philosophical debates other than the debate over naturalism, implications which have gone largely unnoticed to this point. I mention one such debate here.

2.2 *The Instrumentalist Conception and the Ethics of Belief*

Should one believe a proposition for which one lacks evidence if doing so promises to have beneficial consequences? Should one abstain from believing a proposition for which one has a considerable amount of evidence if believing that proposition would have pernicious consequences for oneself or for others? Questions of this sort have been pursued under the rubric 'the ethics

¹⁵ The point is perhaps more obvious in the case of logicism for the following reason. Frege, the father of logicism, was a great *enemy* of empiricism. In contrast, Quine is a great enthusiast for naturalism, and indeed, as we have seen he explicitly suggests something much like the instrumentalist reduction as a way of naturalizing epistemology.

of belief'.¹⁶ My suggestion is that whether the instrumentalist conception is true has crucial implications for the way we should think about such questions. In particular, the truth of the instrumentalist conception is incompatible with certain quite natural positions about the ethics of belief.

Consider, for example, the following very natural reaction to the kind of examples which fuel the ethics of belief literature:

In cases in which what it is epistemically rational to believe clearly diverges from what it is practically advantageous to believe, there is simply no genuine question about what one should believe: Although we can ask what one should believe from the epistemic perspective, and we can ask what one should believe from the practical perspective, there is no third question: what one should believe, all things considered. In any case in which epistemic and practical considerations pull in opposite directions, there is simply nothing to be said about what one should believe all things considered.

Call this view the **Incommensurability Thesis**.

The Incommensurability Thesis is endorsed by Richard Feldman in the course of expressing his skepticism about

the meaningfulness of questions about whether epistemological considerations are outweighed by moral or prudential considerations in figuring out what one ought to do all things considered (Feldman 2000, p. 15).

According to Feldman

Suppose that one belief is prudent for me...but it is not a belief I epistemically ought to have since I lack evidence for it...I can see no values to which we could be appealing when we ask whether the prudential benefit trumps the epistemic cost...There is...no meaningful question about whether epistemic oughts trump or are trumped by other oughts (Feldman 2000, pp. 14-15)¹⁷

However, if the instrumentalist conception of epistemic rationality turns out to be correct, then this would, I think, cast severe doubt upon the Incommensurability Thesis. For if epistemic rationality *just is* instrumental rationality, then there need be no more incommensurability with respect to the ethics of belief than there is within the province of instrumental reason itself.

Consider: among the goals which I hold at the present time are (1) preserving my life and (2) obtaining a chocolate milkshake. The fact that I have these goals gives me reasons to act in certain ways and reasons not to act in other ways. Of course, my holding these two goals might lead to conflicts—conflicts which would not arise if I held either goal in the absence of

¹⁶ For a sampling of the literature, see Heil (1983, 1992), Kelly (2002), Meiland (1980), Mills (1998), Nozick (1993, ch. 3) and Foley (1987, ch. 5).

¹⁷ The possibility that epistemic and practical considerations are incommensurable is raised—but neither endorsed nor discussed at any length—by both Heil (1992, p. 50) and Mills (1998, p. 29).

the other. Suppose, for example, that I can obtain a chocolate milkshake only by engaging in behavior that would place my life in extreme danger. We can imagine a philosopher who insists that, in such circumstances, although we can ask what it is rational for me to do with respect to the goal of obtaining a chocolate milkshake, and we can ask what it is rational for me to do with respect to the goal of preserving my life, there is no third question: what it is rational for me to do all things considered. But this, I think, would not be an impressive suggestion. Because of the way that my goals are ordered with respect to one another, it would be (I can truly report) all-things-considered irrational for me to jeopardize my life in order to obtain a chocolate milkshake.¹⁸

Suppose then that the instrumentalist conception is correct: epistemic rationality is simply instrumental rationality in the service of one's cognitive goals. In that case, it looks as though there will be counterexamples to the Incommensurability Thesis, i.e., cases in which there *is* a fact of the matter about what it is rational to believe all things considered. Suppose, for example, that I can save my life by holding some epistemically irrational belief. Suppose further that the belief concerns some subject matter with respect to which my having true rather than false beliefs is a matter of relative indifference. Now, if epistemic rationality *just is* instrumental rationality, then I think that we can safely conclude: all things considered, it is rational for me to hold this belief, given that I am able to do so. At least, there is no more reason to deny this, than there is to deny that it is instrumentally rational for me to abstain from pursuing a chocolate milkshake in order to save my life. For both cases involve a comparison of the strength of competing *instrumental* reasons.¹⁹

The truth of the instrumentalist conception of epistemic rationality would, I think, undermine the Incommensurability Thesis.²⁰

¹⁸ The present claim should not be misunderstood. I don't mean to commit myself here to the view that it is possible to deliberate rationally about how one's noninstrumental goals or 'final ends' should be ordered. Rather, the point is that *given* the way my goals are in fact ordered, it would be (all-things-considered) irrational for me to jeopardize my life in order to acquire a chocolate milkshake. We might imagine an individual whose preferences are very different from mine; for this person (bizarrely) it is much more important to acquire a chocolate milkshake than to preserve his life. Nothing I have said should be taken as suggesting that it would be all-things-considered irrational for such a person to jeopardize his life in order to acquire the milkshake. (Thanks are due to James Van Cleve for impressing upon me the need to clarify this point.)

¹⁹ Similarly: suppose that, as Chisholm holds, epistemic normativity is really a species of ethical normativity (cf. note 12 above). If so, then in cases in which epistemic considerations and (say) self-interested considerations pull in opposite directions, there need be no more (and no less) incommensurability than there is between ethical and self-interested considerations generally.

²⁰ Foley seems to be well aware that the instrumentalist conception of epistemic rationality has important implications for the ethics of belief. He insists that "All things being considered, it can be rational for an individual to believe what it is not epistemically rational for

But is the instrumentalist conception true? It is to this question which I now turn.

3. Which Cognitive Goals Do We Have?

Perhaps the most serious reason for skepticism about the instrumentalist conception of epistemic rationality is this: what a person has reason to believe does not seem to depend on the content of his or her goals in the way that one would expect if the instrumentalist conception were correct.

It is a characteristic feature of an instrumental reason that one's possessing such a reason is contingent on one's possessing the relevant goal. I have a reason to raise my hand *because* I have the goal of being called upon by the speaker; if I did not have this goal, I would have no such reason. An instrumental reason is a *hypothetical* reason, in the sense that it depends for its existence on the fact that the individual for whom it is a reason possesses a certain goal or goals. This seems to contrast with the categorical character which epistemic reasons apparently possess. On an instrumentalist conception of epistemic rationality, facts about what I have reason to believe are contingent on my possessing certain goals.

One might find this implausible. After all, in our ordinary thought and talk about epistemic reasons, we think and speak of having reasons for belief, not of having reasons for belief insofar as we have goals of such-and-such a sort. We certainly *treat* epistemic reasons as though they are categorical reasons in the course of our ordinary practice. Moreover, we treat epistemic reasons in this way from both the first- and third-person perspectives. That is, one treats epistemic reasons as categorical reasons both in offering such reasons to others as well as in responding to such reasons in the course of one's own theoretical deliberations.

One way of pressing this objection is to appeal to the *intersubjectivity* of epistemic reasons. If both of us know that all of the many previously-observed emeralds have been green, then both of us have a strong reason to believe that the next emerald to be observed will be green, regardless of any differences which might exist in our respective goals. Similarly, in arguing for my conclusions in this paper, I think of myself as attempting to provide strong reasons for believing my conclusions, and *not* as attempting to provide strong reasons for believing my conclusions for those who happen to possess goals of the right sort.

As Tyler Burge notes in a passing remark:

him to believe" (1987, p. 214). He also insists that, although conflicts between epistemic reasons and nonepistemic reasons for belief are sometimes rationally resolvable, it is not the job of a theory of epistemic rationality to resolve them: rather, such questions fall within the jurisdiction of a more general theory of rational belief, a theory which takes into account one's nonepistemic goals (1987, p. 211). These answers, I think, are exactly the answers which an instrumentalist should give to the relevant questions.

Reason has a function in providing guidance to truth, in presenting and promoting truth without regard to individual interest. That is why epistemic reasons are not relativized to a person or to a desire (Burge 1993, p. 475).²¹

There is, I think, a natural response which the instrumentalist might make to this particular line of objection. The instrumentalist might claim that, although in our ordinary practice we treat epistemic reasons as categorical reasons, the relevant aspects of our practice do not constitute evidence for the claim that epistemic reasons are categorical reasons, because these aspects of our practice would be exactly as they are *regardless* of the true nature of epistemic reasons. That is, the reason that we would be inclined to treat epistemic reasons as categorical reasons in the course of our everyday practice, and indeed, to think that epistemic reasons are categorical reasons in the course of our theorizing (regardless of their actual status) is that all of us *do* possess the relevant cognitive goal, viz. believing the truth, or having true rather than false beliefs. Unlike more idiosyncratic goals, which are possessed by some of us but not by others, the goal of believing the truth is a goal which is universally held.²² And if a given goal is sufficiently widespread, it would be quite natural to take that goal for granted in our thought and talk about reasons, and to speak and think, not of reasons for believing relative to that goal, but of reasons for believing *simpliciter*.

Compare: it is natural to think that those of us who have reasons to act in ways which would prolong our lives do so because we have the goal of living longer. Still, it's not surprising that when we present someone with a reason to Φ (where Φ ing is the performing of an action which would lengthen that person's life), we present these reasons as reasons that the individual in question has, and not as reasons that the individual in question has insofar as he or she has the goal of living a longer life. When I see you about to consume a fatally poisonous substance, I might very well think, and say, that you have a reason not to consume the substance. I definitely would *not* think, or say, that you have a reason not to consume the substance insofar as you have the goal of living longer. But these facts about our ordinary practice in no way show that you do have such a reason, independently of your having the relevant goal. For the true story might be this: the goal of living longer is so close to universally-held that we simply take it for granted that any particular person has this goal, and we think and speak accordingly. As Quine might put it: we don't bother to express 'the terminal parameter'. The same might be true with respect to reasons for belief. The apparently categorical character

²¹ Compare Railton (1997, p. 53).

²² "Truth, then, would be rather like what John Rawls has called a primary good, something that is useful for a very wide range of purposes—almost all—and hence will be desired and bring benefit (almost) no matter what our particular purposes might be" (Nozick 1993, p. 68).

of epistemic reasons might actually be an artifact of the universality of the relevant goal.²³

The present dialectical situation should not be misunderstood. One who offers such a story on behalf of the instrumentalist need not claim that the story on offer positively supports the view that epistemic reasons are hypothetical reasons. Rather, the story on offer purports to undermine what would otherwise be extremely strong evidence for the contrary conclusion, viz. that epistemic reasons are categorical reasons. In general, one undermines the claim that *p* is evidence for *q* by showing that *p* would obtain even if *q* was false.²⁴ In the present case, the claim is that the fact that

we constantly think and act as though epistemic reasons are categorical reasons

is evidence for the further claim that

epistemic reasons are categorical reasons.

Let it be conceded that, in general, the fact that we constantly think and act as though such-and-such is the case is strong evidence that such-and-such is the case, all else being equal. In this case though, not all else is equal: what would ordinarily be strong evidence is undermined. Because we would think and act as though epistemic reasons are categorical reasons regardless of their true nature, the fact that we do this does not count as evidence that epistemic reasons are categorical reasons.

The viability of this instrumentalist response, of course, presupposes that there *is* some shared cognitive goal which might underwrite the existence and intersubjectivity of epistemic reasons. In fact, it is here, I believe, where the instrumentalist conception of epistemic rationality founders: there is simply no cognitive goal or goals, which it is plausible to attribute to people generally, which is sufficient to account for the relevant phenomena. Individuals do not typically have this goal: believing the truth.

The sense in which individuals typically lack this goal requires clarification. No doubt, individuals frequently manifest a preference for having true beliefs about particular subject matters. Thus, individuals seek out reliable

²³ Compare Kant on the pervasiveness of happiness as an end. For Kant, our reasons to perform actions conducive to our own happiness have hypothetical force as opposed to the categorical force of moral reasons. But because we all have the end of happiness as a matter of 'natural necessity', we state imperatives of prudence, like imperatives of morality, in 'assertoric' rather than 'hypothetical' form (1981, p. 26).

²⁴ You suggest that the fact that my dog is barking is evidence that she wants to go outside, I undermine this claim by informing you that my dog barks constantly, regardless of whether she wants to go outside. Cf. Pollock's excellent discussion of epistemic defeasibility (Pollock, pp. 37-39). In Pollock's terminology, we are concerned here with 'undercutting' as opposed to 'rebutting' defeaters.

sources in order to ask for directions about how to arrive at a particular destination, look up facts in books, visit museums, read newspapers, and watch news programs in order to acquire accurate information. Individuals perform scientific experiments and conduct statistical surveys. All of these activities, I think, are indicative of a concern for truth. Even an action as simple as redirecting one's gaze from the center of the room to the corner in order to discover the cause of an unexpected sound is (perhaps) indicative of a concern for truth.

But activities such as these indicate only that the individual in question has fairly specific, particularized cognitive goals. When I ask a reliable source for directions to Fenway Park, I do so because it is important to me to have true beliefs about how to get to Fenway Park. (About this subject matter I have a strong preference for having true beliefs rather than false beliefs, and for having true beliefs to no beliefs at all.) Similarly, when, upon hearing a strange noise in the corner of the room, I intentionally redirect my gaze in order to discover its source, this behavior is indicative of the fact that I have a quite specific cognitive goal: that of finding out (the truth about) what's happening in the corner of the room. Parallel remarks apply to the cases of scientific experiments and statistical surveys.

Of course, some cognitive goals are *wider* than others. When I consult a reliable source in order to acquire accurate information about how to get to Fenway Park, I have one particular question to which I want a true answer: "How do I get to Fenway Park?"²⁵ My goal of believing the truth about how to get to Fenway Park is a relatively *narrow* goal, in the following sense: there is a fairly limited range of information which is such that, if I came into cognitive possession of this information, my doing so would constitute this goal's being better achieved. On the other hand, when I read the morning newspaper or watch a television news program, there is (typically) not some one question or small range of questions which I want answered. Rather, I am typically motivated to undertake such activities because I have the goal of, e.g., acquiring information about any event of significance which has recently occurred. The goal which motivates my reading the newspaper is a relatively *wide* goal, in the sense that there are *many* truths (a fairly wide range of information) such that my coming to believe (any of) these truths would constitute the relevant goal's being better achieved.

There are, however, very real limits to how wide even the widest of my cognitive goals are. In addition to those many truths such that my believing them would contribute to the achievement of some goal that I have, there are also (countless) truths such that my believing them would not contribute to

²⁵ Of course, the fact that I have the goal of finding out how to get to Fenway Park will often give rise to other goals: if I am told that in order to get to Fenway Park, I first have to get to point X, I will acquire the goal of finding out how to get to point X, etc.

any goal that I actually have. Whether Bertrand Russell was right- or left-handed, whether Hubert Humphrey was an only child—these are matters of complete indifference to me. That is, I have no preference for having true beliefs to having no beliefs about these subjects; nor, for that matter, do I have any preference for having true beliefs to false beliefs. There is simply no goal—cognitive or otherwise—which I actually have, which would be better achieved in virtue of my believing true propositions about such subjects, or which would be worse achieved in virtue of my believing false propositions about them.

However, from the fact that some subjects are matters of complete indifference to me, it does not follow that I will inevitably lack epistemic reasons for holding beliefs about those subjects. If, despite my utter lack of interest in the question of whether Bertrand Russell was left-handed, I stumble upon strong evidence that he was, then I have strong epistemic reasons to believe that Bertrand Russell was left-handed. Indeed, my epistemic reasons will be no different than they would be if I had acquired the same evidence deliberately, because I *did* have the goal of finding out whether Russell was left-handed. Once I come into possession of evidence which strongly supports that claim that *p*, then I have epistemic reasons to believe that *p*, regardless of whether I presently have or previously had the goal of believing the truth about *p*, or any wider goal which would be better achieved in virtue of my believing the truth about *p*. The fact that I can have epistemic reasons to believe propositions even though doing so holds no promise of better achieving any of my goals (cognitive or otherwise) fits poorly with the instrumentalist conception of epistemic rationality, since whether it is instrumentally rational to Φ always depends on the contents of one's goals.²⁶

It is for this reason that the instrumentalist conception of epistemic rationality fails to do justice to the intersubjectivity of epistemic reasons. For

²⁶ Harman (1999b) is similarly skeptical of the idea that individuals typically possess 'a general desire' to 'believe what is true and not believe what is false'. He writes: 'Of course, people do not actually have this general desire. Curiosity is more specialized. One wants to know whether *P*, who did *D*, what things are *F*, and so forth' (p. 100).

I strongly agree with Harman's claim that people do not have the general desire in question. However, it would be a mistake, I think, to assimilate our curiosity, or our concern with truth, to the desire to know the answers to specific questions (as Harman seems to suggest here). When I read the morning newspaper, I am *sometimes* motivated to do so because I have the goal of discovering the answer to some specific question (e.g., who won last night's election). More frequently, however, I am motivated to do so not because I want to find the answer to any particular question; rather, I simply want to learn interesting and important truths about the world.

In attempting to characterize our concern with truth, there are two opposite errors that must be avoided:

- (i) that our concern with truth is such that it is better satisfied whenever we come to believe any true proposition, no matter how trivial or insignificant, and
- (ii) that our concern with truth is wholly exhausted by our wanting to know the answers to specific questions.

individuals will typically differ greatly with respect to which subject matters are matters of indifference and which are not. That is, individuals will differ greatly with respect to *which* cognitive goals they possess. Among my cognitive goals is the goal of having true rather than false beliefs about the nature of epistemic rationality. But this is no doubt an extremely idiosyncratic goal relative to the general population: very few people, I suspect, have some goal which would be better promoted in virtue of having true beliefs about the nature of epistemic rationality. Because I live in Somerville, Massachusetts, I have a strong interest in having true rather than false beliefs about which Somerville streets are one-way streets; because I do not live in Bakersfield, California (and have no intention of going there) I have no interest in having true beliefs about which Bakersfield streets are one-way. Someone who lives in Bakersfield is likely to differ from me in both of these respects. It does not follow that we will inevitably differ in what we have epistemic reason to believe. Differences in our cognitive goals need not find reflection in the epistemic reasons that we possess.

Not only are there (many) subjects with respect to which I have no preference for having true beliefs, there are also subjects with respect to which I would prefer to have no beliefs at all to having true beliefs. Thus, I tend to see newly-released movies after many of my friends. During the interval of time which is bounded on one side by my friends' viewing of the movie and bounded on the other side by my viewing the movie, I often make a conscious, deliberate effort to avoid finding out how the movie ends—since doing so might very well interfere with my enjoyment when I do see it. (When conversations about the movie begin in my presence, I either excuse myself or, reminding the discussants that I have yet to see the movie, implore them not to “give away” the ending, and so on.) That is, I quite deliberately take steps to avoid acquiring information about the movie. Sometimes these efforts are successful, sometimes they are not. When they are unsuccessful—as when someone inconsiderately blurts out the ending in my presence—it does not follow that I have no epistemic reasons to believe the propositions which he asserts. Indeed, with respect to the question of which epistemic reasons I possess, there is no difference between this case and a case in which I *ask* the individual to tell me the ending because I *do* have some goal which would be better achieved by my believing the relevant truths. The fact that in the one case I do have a goal which is better achieved by my believing the relevant truths, but in the other I have no such goal—indeed, I have goals which would be hindered or frustrated by my believing these truths—makes no difference to my epistemic reasons.

In his brief consideration of the putative possibility that someone might lack the goal of now having true beliefs and not now having false beliefs (1987, pp. 11-12), Foley speculates that this envisaged possibility might

turn out not to be possible after all. In support of this speculation, Foley claims that “the vast majority of us” attach an intrinsic value to having true beliefs, and even those among us who do not do so presumably care about having true beliefs because they recognize that such beliefs have instrumental value. Now, perhaps it is in fact impossible (in some fairly weak sense of “impossible”) for someone to be *wholly unconcerned* with having true beliefs, in the sense that, necessarily, every individual is such that there are some subject matters about which he or she is concerned to believe the truth. Perhaps it is even the case that, as Ernest Sosa has suggested, “for any arbitrary belief we actually hold, we would prefer that it be true rather than not be true, other things equal”.²⁷ But this—as Sosa himself notes—is a far cry from the claim that individuals typically have some goal which is better achieved *whenever* one believes some true proposition, no matter how trivial or insignificant. But of course, one can have extremely strong epistemic reasons to believe utterly trivial and insignificant propositions.

Ultimately, Foley appears prepared to say that, if a person genuinely *did* lack the requisite goal—which he somewhat grudgingly admits *may* be possible—then *nothing* would be either epistemically rational or irrational for that person (1987, p. 12). A similar conclusion is embraced by David Papineau.²⁸ Interestingly, Papineau takes the possibility of individuals who lack the requisite cognitive goals as *favoring* the kind of instrumentalist account which I am attempting to undermine. He argues as follows. After noting the existence of cases in which individuals deliberately avoid seeking evidence in order to avoid unwanted beliefs, he claims (correctly, I believe) that there are cases of this sort in which the individuals in question are subject to no legitimate criticism for acting in this manner. He then concludes that this supports the idea that epistemic norms—or “norms of judgement” have a “hypothetical” as opposed to a “categorical” character. But to proceed in this way is to conflate (1) the reasons which one may or may not have to seek out further evidence which bears on the truth of *p*, and (2) the reasons which one may or may not have to believe *p*.²⁹

²⁷ Sosa (2002, p. 3).

²⁸ Papineau, “Normativity and Judgement”. See especially pages 23-25.

²⁹ In a passing footnote (p. 24, fnt.8), Papineau shows that he is aware of the distinction. He seems, however, not to appreciate its potential significance, for in the main text he passes, directly and without argument, from the claim that (i) an individual might be under no obligation to gather evidence which she does not presently possess to the claim that (ii) an individual might be under no obligation to conform her beliefs to the evidence which she presently possesses. In the same footnote, Papineau notes that deliberately refusing to conform one’s beliefs to evidence which one already possesses (in contrast to deliberately refusing to seek out further evidence) is of “doubtful psychological possibility” and wonders why this is so. Below, I will suggest that the asymmetry in question gives us further reason to doubt the sort of view which Papineau favors.

Are there any positive reasons for supposing, against Foley and Papineau, that an individual might have reasons to hold beliefs about a subject matter even if she has no goal which would be better promoted in virtue of her believing the relevant truths? Consider the following. When I undertake deliberate measures in order to avoid discovering how the movie ends, my project is simply this: I want to avoid the acquisition of reasons for believing the truth about how the movie ends. Notice, however, that if the possibility of acquiring reasons for believing the truth about *p* is contingent on one's having some goal which would be better promoted by believing the truth about *p*, then this project is incoherent: there is no need to deliberately avoid the acquisition of epistemic reasons to believe propositions about subjects with respect to which one has no desire to believe the truth, for one knows *a priori* that there are no such reasons. (Indeed, that there *could not be* such reasons.) But in fact, the envisaged project is not incoherent. I might have epistemic reasons to believe the truth about how the movie ends despite my not having the relevant goal, as becomes apparent when—in spite of my best efforts—I acquire the unwanted belief by stumbling upon the unwanted reasons. Notice that when I acquire the unwanted belief in this fashion, that I do so is not merely a matter of pure psychological compulsion: in such circumstances, we might very well *explain why* I formed the unwanted belief by *citing* my epistemic rationality, along with the fact that I was presented with epistemic reasons of the relevant sort. But this explanation would not be available to us if we claimed, with Foley, that nothing would be epistemically rational for one who lacked the relevant goal. Put simply: one cannot *immunize* oneself against the possibility of acquiring reasons for belief by not caring about the relevant subject matter.

Philosophers often suggest that in addition to our many and various *local* cognitive aims (e.g., having true rather than false beliefs about what the weather will be like tomorrow, or about whether the stock market will continue to go up), there is some more general, *global* cognitive aim with respect to which our epistemic practices and efforts are to be assessed. Thus, Roderick Chisholm once suggested that we should understand our central cognitive goal as that of 'having the largest possible set of logically independent beliefs that is such that the true beliefs outnumber the false beliefs'.³⁰ As we have seen, Foley suggests that judgements of epistemic rationality are made relative to the goal of 'now believing true propositions and not now believing false propositions'. Here, of course, there is not one cognitive goal but rather two: (1) now believing true propositions and (2) not now believing false propositions. This fact leads directly to familiar questions

³⁰ Chisholm (1982, p. 7). It is, I think, extremely dubious that anyone has ever had, or should have had, this particular goal. (Notice that the goal to which Chisholm refers would be better achieved by someone who has 5,000,000 true beliefs and 4,999,999 false beliefs than by someone who has 9,999,998 true beliefs and zero false beliefs.)

about the relative weights which are to be assigned to these two goals.³¹ Typically, questions about how the central cognitive aim is to be understood are raised only to be set aside, as by Alston:

Epistemic evaluation is undertaken from what we might call the 'epistemic point of view'. That point of view is defined by the aim of maximizing truth and minimizing falsity in a large body of beliefs. The qualification 'in a large body of beliefs' is needed because otherwise one could best achieve the aim by restricting one's beliefs to those that are obviously true. That is a rough formulation. How large a body of beliefs should we aim at? Is any body of beliefs of a given size, with the same truth-falsity ration equally desirable? And what relative weights should be assigned to the two aims of maximizing truth and minimizing falsity? We can't go into all that here; in any event, however these issues are settled, it remains true that our central cognitive aim is to amass a large body of beliefs with a favorable truth-falsity ratio.³²

Here, the suggestion seems to be that it would be a good thing if we had answers to all of these questions; that these issues might be settled if only we had the time and space to go into them (indeed, that they *will* be settled at some point), and that their being definitively settled would be a desirable state of affairs.

But there is no reason to think that we do possess some one "central cognitive aim" in the relevant sense. That is, there is no aim, or goal, which (1) is better achieved whenever one adds true propositions or avoids adding false propositions to one's stock of beliefs, and which (2) people actually hold.³³ At least, nothing in the way that people behave suggests that they do have such a central cognitive goal—as opposed to a vast number of more special-

³¹ As is often noted, there is a certain tension between pursuing these two cognitive goals. The more weight one gives to the goal of not believing false propositions, the more it behooves one to be very *conservative* about what one believes, and to believe only those propositions for which one has a *great deal* of evidence. On the other hand, the more weight one gives to the believing of true propositions, the more it behooves one to believe large numbers of propositions, including propositions for which one does not have a great deal of evidence. Epistemic caution and abstinence count as virtues relative to the aim of not believing false propositions; epistemic aggressiveness and commitment count as virtues relative to the aim of believing true propositions.

The point that there is a certain tension between pursuing the two cognitive goals was first made (I believe) by William James in his (1956, pp. 18-19).

³² Alston (1989, pages 83-84). Even the little that Alston says here is enough to raise qualms. For example, Alston's remark that

the qualification 'in a large body of beliefs' is needed because otherwise one could best achieve the aim by restricting one's beliefs to those that are obviously true

seems to suggest that the relevant qualification is motivated by the thought that there is *not* a large number of obvious truths. But, given any natural interpretation of 'a large number' and 'obvious', this thought is mistaken.

³³ Here I am tempted to make the strong claim that no one, or almost no one, actually holds such a goal. In fact, given the intersubjectivity of epistemic reasons, it is enough for my argument, I think, if there is even a single person who lacks such a goal. I believe that I am such a person.

ized, narrower cognitive goals. After all, people routinely pass up opportunities to add true beliefs to their present stock even when doing so would be of little or no cost. Nor, I think, does introspection reveal the existence of any such goal. (Here I speak for myself, and invite the reader to undertake a similar inquiry.)

There is, I suspect, a very good reason why the question of the relative weights which are to be given to the two cognitive goals—believing what is true and not believing what is false—is typically raised only to be set aside. Quite simply: there is no answer to this question when it is asked at the level of extreme generality at which it is typically posed. As statisticians are fond of emphasizing, the relative importance of avoiding a ‘Type I mistake’ (that is, failing to take something to be true which is in fact true) as opposed to avoiding a ‘Type II mistake’ (that is, taking something to be true which is in fact false) is highly sensitive to specific features of a given context.

When it is instrumentally rational for me to Φ , this is because Φ ing promises to promote some goal or goals which I possess. The attempt to assimilate epistemic rationality to instrumental rationality founders on the fact that one can have epistemic reasons to believe propositions even in cases in which it is clear that one’s believing those propositions holds no promise of advancing any goal which one actually possesses.

4. Some Instrumentalist Replies

In this section, I take up some natural instrumentalist replies to the preceding argument. Rather than engaging in a futile effort to consider every reply which an instrumentalist might offer, I want to examine, at some length, what I take to be the two most formidable and philosophically interesting replies.

I have argued that the attempt to assimilate epistemic rationality to instrumental rationality founders on the fact that it can be epistemically rational to believe propositions even in cases in which it is clear that believing those propositions would not advance any goal which one actually holds. The first instrumentalist reply challenges the claim that one has no goal which is better promoted whenever one believes a proposition that one has epistemic reasons to believe. According to this reply, individuals generally do have such a goal, but I have missed this fact as a result of unduly restricting the kinds of considerations which can justify attributing goals to individuals. The second instrumentalist reply grants the claim that individuals typically do not have such goals, but contends that this fact is compatible with the instrumentalist conception of epistemic rationality.

4.1 Truth as The Constitutive Aim of Belief

Recall our earlier example: despite the fact that I would prefer not to believe the truth about how the movie ends, I can acquire epistemic reasons to believe the truth about how the movie ends.

But *don't* I have the goal of believing the truth about how the movie ends? You inconsiderately blurt out the ending of the movie in my presence ('the butler did it'); in response, I immediately come to believe that the butler did it. I now have a belief about how the movie ends—and beliefs, as Bernard Williams has famously claimed, "aim at truth".³⁴ Truth is the *constitutive aim* of belief. Perhaps then, in virtue of my newly-acquired status as one who has beliefs about how the movie ends, I *inherit* the aim or goal of believing the truth about how the movie ends, in virtue of the nature of belief. Beliefs aim at truth; I am a believer about x; therefore, I have the aim of believing the truth about x.

This line of thought, I believe, is fallacious. After all, why is the argument

I have beliefs about x
The aim of any belief is truth
Therefore, I have the aim of having beliefs about x which are true

any better than the following (presumably bad) argument?

I have a heart
The aim of any heart is to pump blood
Therefore, I have the aim of having a heart which pumps blood.

(Someone is attempting to commit suicide by stopping his own heart from pumping blood: It would be a mistake, I take it, to attribute to such a person the aim of having a heart which pumps blood.)

Moreover, the crucial premise—that "belief aims at truth"—is notoriously obscure. Talk of belief "aiming" at truth is, I assume, metaphorical, and this metaphor has yet to be fully unpacked.³⁵ Suppose, however, that there is some non-metaphorical interpretation of "belief aims at truth" or "truth is the constitutive aim of belief" which is both true and philosophically interesting. Would this help the instrumentalist?

In fact, to appeal to the claim that belief aims at truth at this juncture is, I think, essentially to abandon the attempt to reduce epistemic rationality to instrumental rationality. After all, what is distinctive about instrumental

³⁴ Williams (1973). The general line of thought is well-summarized by one of its critics, Fred Dretske (Dretske, 2000).

³⁵ Although much progress toward this end has undoubtedly been made by Velleman (2000b) and Wedgwood (2002).

rationality is precisely the fact that which instrumental reasons a person has depends on which ends (goals, aims) he or should would prefer to have realized. To appeal to states which by their very nature ought to be a certain way (regardless of whether anyone has any preference for their being that way) is already to move beyond *instrumental* rationality.

Compare: a neo-Aristotelian might hold that, because human beings are by their very nature certain sorts of beings, they have distinctive ends which they ought to realize (regardless of whether they have any preferences for the realization of those ends) and the fact that they have these ends gives them reasons to act in some ways rather than in others. Such reasons, I think, would not be instrumental reasons.

I want to ask the speaker a question, and I know that I will only be able to ask my question if I raise my hand. These facts give me an instrumental reason to raise my hand. In explaining why my action is rational, there is no need to appeal to a "constitutive aim" of my action, or some such thing. That is, no role is played by constitutive aims in paradigmatic exhibitions of instrumental rationality.

4.2 *The Appeal to Merely Hypothetical Goals*

Consider the following instrumentalist response:

Epistemic rationality is in fact simply a special case of instrumental rationality (viz., instrumental rationality in the service of some cognitive goal), but it is not crucial that individuals actually do possess the relevant goal. Rather, as theorists we can evaluate how well an individual's ways of revising his or her beliefs *would* promote the goal in question, regardless of whether he or she in fact possesses that goal. And it is from this perspective that judgements of 'epistemically rational' or 'epistemically irrational' are made.

Imagine a being who differs from us only in that he is afflicted with a peculiar sort of avarice: he always strongly prefers to believe more truths (no matter how trivial or useless for his *other* projects) to fewer. And he loathes the thought of believing anything false. Plausibly, the most instrumentally rational strategy for such a person to pursue is to believe all and only those propositions which it is epistemically rational for him to believe. Does this fact help the instrumentalist?

We can, of course, consider how it would be rational for an individual to pursue some goal whether or not the individual actually holds that goal. Thus, we can ask how it would be rational for me to pursue the goal of now believing true propositions and now not believing false propositions, even though I don't in fact have this goal.³⁶ But this, I think, is not enough to save the instrumentalist conception of epistemic rationality. The crucial fact here is the following: whether it is in fact instrumentally rational for me to

³⁶ As noted by Foley (1987, p. 12).

Φ depends on the content of the goals which I *actually* hold. It's no doubt true that *if* I had the goal of asking the speaker a question, I would have an instrumental reason to raise my hand, and (all else being equal) my doing so would be instrumentally rational. But if in fact I do not have this goal, I have no reason to raise my hand. Only goals which I actually hold make a difference to what is instrumentally rational for me.³⁷ But I can have epistemic reasons to believe propositions even though doing so holds no promise for promoting any goal which I actually hold. This suggests that it is a mistake to assimilate epistemic rationality to instrumental rationality.

No doubt, much of the allure of the instrumentalist conception consists in the fact that

- (1) Insofar as I am pursuing the goal of now believing true propositions and not now believing false propositions, it is a good (i.e., instrumentally rational) strategy to (i) believe those propositions which it is epistemically rational for me to believe and (ii) to not believe those propositions which it is epistemically irrational for me to believe.

I believe that (1) is true. But the correctness of the instrumentalist conception of epistemic rationality cannot be derived from (1).

In order to appreciate this fact, consider a parallel case. It's no doubt true that

- (2) Insofar as I am pursuing the goal of being a moral person, it is a good (i.e., instrumentally rational strategy) for me to (i) perform all of those actions which I have overriding moral reasons to perform, and (ii) to not perform any action which I have overriding moral reasons not to perform.

But no one would think that it follows from (2) that morality *just is* instrumental rationality, or some such thing. After all, even a Platonist about moral reasons would presumably accept (2). Similarly, from the fact that it is a good (i.e., instrumentally rational) strategy to be epistemically rational insofar as one is pursuing the goal of now believing true propositions and

³⁷ Perhaps Williams (1981) can be interpreted as arguing for a somewhat wider conception of instrumental rationality, according to which I can have instrumental reasons to advance not only goals which I actually hold but also goals which I might reach by a process of sound deliberation from my present 'subjective motivational set'. If this is in fact the right way of thinking about instrumental rationality, then the objection which I have developed should be put like this: it can be epistemically rational to believe propositions even when it is clear that doing so would promote no goal which one actually holds or which might be reached by a process of sound deliberation from one's subjective motivational set.

not now believing false propositions, it in no way follows that epistemic rationality just is instrumental rationality in the service of this goal.

5. The Role of Instrumental Rationality in Theoretical Reasoning: Theoretical Rationality as a Hybrid Virtue

I have argued that it is a mistake to attempt to assimilate epistemic rationality to instrumental rationality in the service of one's cognitive goals. Nevertheless, it would also be a mistake to underestimate the epistemic importance of instrumental rationality. Indeed, responding to instrumental reasons plays a pervasive and indispensable role in both theoretical inquiry and theoretical reasoning—a role which complements the role of epistemic rationality. In the final section of this paper, I want to delineate, in broad outline, the respective roles of epistemic and instrumental rationality in the achievement of one's cognitive goals through theoretical reasoning.

Let's begin by considering the role of instrumental rationality in the activities constitutive of theoretical inquiry. As we have already noted, the fact that one has certain cognitive goals often makes it instrumentally rational for one to act in some ways rather than others. In particular, the fact that one has the goal of finding out the truth about some question often provides an instrumental reason to *improve one's epistemic position* with respect to that question. Suppose that I hear a strange and unexpected sound behind me, and, seeking to find out the source of this noise, I turn around. Here, the reason that I have to turn around is an instrumental reason—I have the (cognitive) goal of finding out what is responsible for the relevant noise, and given this goal, it is instrumentally rational for me to change my epistemic position in a certain way.³⁸ Suppose further that, upon turning around, I discover the source of the noise: a cat has entered the otherwise-empty room. Finding myself face-to-face with the cat, it is now epistemically rational for me to believe that a cat was responsible for the noise. What is the relationship between my possessing an epistemic reason to believe this proposition and my possessing the relevant cognitive goal? In one respect, the fact that it is epistemically rational to believe this proposition does not depend on the fact that I possess the goal: someone who occupied my same epistemic position, but who lacked the goal, would have the same epistemic reason that I do. On the other hand, that it is epistemically rational for me to believe that a cat is responsible for the noise is historically dependent on my possession of the relevant goal: if I did not possess the relevant goal, I would never have turned around and (hence) never acquired epistemic reasons to believe the proposition. Notice that, in this case, fulfilling my goal of dis-

³⁸ In this case, changing my epistemic position in the requisite way involves changing my physical position, but, as will become clear below, deliberately changing one's epistemic position does not always involve changing one's physical position.

covering the truth about the source of the noise requires that I exhibit *both* instrumental rationality and epistemic rationality: it is because I am instrumentally rational that I improve my epistemic position in the requisite way, and it is because I am epistemically rational that, having improved my epistemic position, I come to the true belief that a cat is responsible for the noise.

At the most abstract level, scientific inquiry itself might be understood as simply a (much) more complicated and sophisticated version of this basic picture. The reasons which one has to engage in practices of evidence-gathering and experimentation are instrumental reasons; once the experiments have been performed, however, what it is rational to believe is no longer a matter of instrumental (but rather epistemic) rationality.

Moreover, being instrumentally rational in the pursuit of one's cognitive goals plays an important role in theoretical reasoning itself. By *theoretical reasoning* I mean reasoning which is undertaken in order to determine what to believe (as opposed to practical reasoning, reasoning which is undertaken in order to determine what to do). Theoretical reasoning, I believe, closely resembles theoretical inquiry in that the former, like the latter, involves responding to both epistemic and instrumental reasons.

The capacity to respond to instrumental reasons is central to theoretical reasoning because of the *directed* or *goal-oriented* nature of such reasoning. In reasoning theoretically, one does not simply arrive at new beliefs by applying rules of inference willy-nilly to one's present corpus. Rather, in engaging in theoretical reasoning, one typically has some particular question or questions which one wants answered. That is, one has a certain cognitive goal which one wants to achieve, and the content of this goal gives one instrumental reasons to engage in certain mental activities rather than others.

Consider, for example, the activity of *calculating*. At the conclusion of meals, I am often confronted with the task of determining how much to leave as a gratuity, given that I want to leave an amount which is equal to 20% of the total bill. Typically, I pursue the relevant cognitive goal by first, determining how much 10% of the total bill would be and then doubling that number. Even a process of reasoning as simple as this, I believe, involves responding to both epistemic and instrumental reasons. Before beginning my calculation, I don't know how much to leave, and I pursue the goal of determining how much to leave by attempting to improve my epistemic position with respect to the relevant question. As we saw above, the reasons which one has to improve one's epistemic position with respect to some question are typically instrumental reasons—although in this case, responding to such reasons does not involve changing my physical position. Rather, I respond to these instrumental reasons by undertaking those mental activities which I need to undertake in order to arrive at a solution. Thus, the rationality which I

exhibit in undertaking the sub-task of determining 10% of the total bill is instrumental rationality in the service of my cognitive goal (if I did not have the cognitive goal of determining 20% of the bill I would quite literally have *no reason* to undertake this task). But having performed any particular step in the calculation, that I believe what I should believe given my newly-arrived at epistemic position is a matter of my being epistemically rational, i.e., appropriately sensitive to the epistemic position which I have now come to occupy. My epistemically rational belief that 10% of the total bill is n can then be used as an input to further reasoning, reasoning which it is (instrumentally) rational for me to undertake in virtue of my particular cognitive goal. In this way, instrumental rationality and epistemic rationality work in tandem in cases in which an individual pursues his or her cognitive goals through theoretical reasoning.

The general point, viz. that which cognitive goals one possesses can and should make a difference to which conclusions one ultimately reaches through theoretical reasoning, is noted by Harman, who provides the following example:

There are various conclusions that Jack could reach right now...He could solve some arithmetical problems...He could try to resolve a philosophical paradox...But, at the moment, Jack is locked out of his house and really ought to try to figure out where he left his keys. If Jack thinks about where he left his keys, however, he won't be able at the same time to resolve the philosophical paradox or solve the arithmetical puzzles. Because he wants very much to get into his house, he devotes his attention to figuring out where his keys must be.

From this, Harman concludes that "your desires can rationally affect your theoretical conclusions by affecting what questions you use theoretical reasoning to answer" (Harman 1999a, p. 15).

All of this, I think, is correct. However, what I would like to emphasize is that practical, goal-oriented considerations enter in not only at the most general level of, say, deciding to determine the present location of one's keys as opposed to spending time attempting to solve a philosophical paradox. Rather, even *after* one has adopted the goal of determining the present location of one's keys, or the goal of determining how much one should leave as a gratuity, the subsequent pursuit of such adopted goals *via* theoretical reasoning will typically require responding to instrumental as well as epistemic reasons. Theoretical reasoning of any significant degree of complexity requires responsiveness to both epistemic and instrumental reasons.

Philosophers tend to juxtapose theoretical rationality and practical rationality, where being practically rational consists in being responsive to practical reasons. This suggests (naturally enough) that theoretical rationality similarly consists in being responsive to a certain *type* of reason, viz. 'theoretical reasons'. 'Theoretical rationality' is thus sometimes used as a synonym for 'epistemic rationality', and 'theoretical reason' as a synonym for 'epistemic

reason'.³⁹ But this terminology, I think, is both symptomatic of, and further encourages, a mistaken view about the nature of theoretical rationality. Theoretical rationality is a virtue which consists in proficiency in theoretical reasoning. Being proficient in theoretical reasoning, in turn, involves manifesting sensitivity to two different kinds of reasons: epistemic reasons and those instrumental reasons which one possesses in virtue of possessing the particular cognitive goals which one does in fact possess. We might imagine—with some difficulty, perhaps—a person who has either sensitivity in the absence of the other. That is, we can imagine a being who is perfectly epistemically rational (in the sense that at any given moment she believes all and only those propositions which it is epistemically rational for her to believe at that time) but who constantly fails to undertake those mental activities which she needs to undertake in order to achieve her cognitive goals. On the other hand, we can imagine a being who, being fully instrumentally rational, does undertake the needed mental activities but fails to achieve his cognitive goals in virtue of being pathologically epistemically irrational. Both of these two individuals should, I think, be considered seriously deficient with respect to their possession of the virtue of theoretical rationality.

There is not some one kind of reason—'theoretical reasons'—sensitivity to which qualifies one as theoretically rational. Rather, being theoretically rational is a *hybrid* virtue: it involves sensitivity to two very different kinds of reasons.⁴⁰

References

- Alston, William 1989: "Concepts of Epistemic Justification", reprinted in his *Epistemic Justification* (Ithaca: Cornell University Press), pp. 81-113.
- Audi, Robert 1991: *Practical Reasoning* (London: Routledge).
- Burge, Tyler 1993: "Content Preservation", *Philosophical Review* 103, pp. 457-88.
- Chisholm, Roderick 1957: *Perceiving*. (Ithaca: Cornell University Press).
- Chisholm, Roderick 1982: *The Foundations of Knowing*. (Minneapolis: University of Minnesota Press).
- Chisholm, Roderick 1991: "Firth and the Ethics of Belief", *Philosophy and Phenomenological Research* 51, pp. 119-127.

³⁹ See for example Audi (1991).

⁴⁰ For helpful comments on earlier drafts of this paper I am grateful to Robert Nozick, Derek Parfit, Jim Pryor, Richard Heck, Pamela Hieronymi, and Aaron James. For financial support, I am grateful to the Society of Fellows at Harvard University.

- David, Marian 2001: "Truth as the Epistemic Goal", in Mathias Steup (ed.) *Knowledge, Truth, and Duty*. (Oxford: Oxford University Press), pp. 151-167.
- Dretske, Fred 2000: "Norms, History, and the Constitution of the Mental", reprinted in his *Perception, Knowledge and Belief*. (Cambridge: Cambridge University Press), pp. 242-258.
- Feldman, Richard 2000: "The Ethics of Belief", *Philosophy and Phenomenological Research* LX, pp. 667-695.
- Field, Hartry 2000: "A Prioricity as an Evaluative Notion", in Boghossian and Peacocke (ed.) *New Essays on the A Priori*. (Oxford: Oxford University Press).
- Firth, Roderick 1998a: "Chisholm and the Ethics of Belief" reprinted in Troyer (ed.) *In Defense of Radical Empiricism: Essays and Lectures by Roderick Firth*. (New York: Rowman and Littlefield), pp. 143-155.
- Firth, Roderick 1998b: "Are Epistemic Concepts Reducible to Ethical Concepts", reprinted in Troyer (ed.) *In Defense of Radical Empiricism: Essays and Lectures by Roderick Firth*. (New York: Rowman and Littlefield), pp. 237-249.
- Foley, Richard 1987: *The Theory of Epistemic Rationality*. (Cambridge: Harvard University Press).
- Foley, Richard 1993: *Working Without A Net*. (Oxford: Oxford University Press).
- Giere, Ronald 1988: *Explaining Science: A Cognitive Approach*. (Chicago: University of Chicago Press).
- Giere, Ronald 1989: "Scientific Rationality as Instrumental Rationality", *Studies in History and Philosophy of Science* 20, pp. 377-384.
- Goldman, Alvin 1986: *Epistemology and Cognition*. (Harvard University Press: Cambridge).
- Hampton, Jean 1998: *The Authority of Reason* (Cambridge: Cambridge University Press).
- Harman, Gilbert 1999a: "Rationality", reprinted in his *Reasoning, Meaning, and Mind* (Oxford: Clarendon Press), pp. 9-45.
- Harman, Gilbert 1999b: "Pragmatism and Reasons for Belief", reprinted in his *Reasoning, Meaning, and Mind* (Oxford: Clarendon Press), pp. 93-116.
- Heil, John 1983: "Believing What One Ought", *Journal of Philosophy* 80, pp. 752-765.
- Heil, John 1992: "Believing Reasonably", *Noûs* 26, pp. 47-62.
- Jackson, Frank 1999: "Non-Cognitivism, Normativity, and Belief", *Ratio* XII, pp. 420-435.
- James, William 1956: "The Will to Believe", in his *The Will to Believe and Other Essays in Popular Philosophy* (New York: Dover), pp. 1-31.

- Kant, Immanuel 1981: *Grounding for the Metaphysics of Morals*. (Indianapolis: Hackett).
- Kelly, Thomas 2002: "The Rationality of Belief and Some Other Propositional Attitudes", *Philosophical Studies* 110, pp. 163-196.
- Kim, Jaegwon 1993: "What is 'Naturalized Epistemology'?" reprinted in his *Supervenience and Mind* (Cambridge: Cambridge University Press), pp. 216-236.
- Kitcher, Philip 1992: "The Naturalist's Return", *The Philosophical Review* 101, pp. 53-114.
- Kornblith, Hilary 1993: "Epistemic Normativity", *Synthese* 94, pp. 357-376.
- Korsgaard, Christine 1997: "The Normativity of Instrumental Reason" in Garrett Cullity and Berys Gaut (eds.) *Ethics and Practical Reason* (Oxford: Clarendon Press), pp. 215-254.
- Laudan, Larry 1990a: "Normative Naturalism", *Philosophy of Science* 57, pp. 44-59.
- Laudan, Larry 1990b: "Aimless Epistemology?", *Studies in History and Philosophy of Science* 21 (1990b), pp. 315-322.
- Laudan, Larry 1996: *Beyond Positivism and Relativism* (Boulder, Colorado: Westview Press).
- Mackie, J.L. 1979: *Ethics: Inventing Right and Wrong* (Harmondsworth: Penguin).
- Maffie, James 1990a: "Recent Work on Naturalized Epistemology", *American Philosophical Quarterly* 27, pp. 281-293.
- Maffie, James 1990b: "Naturalism and the Normativity of Epistemology", *Philosophical Studies* 59, pp. 333-349.
- Meiland, Jack 1980: "What Ought We to Believe? or The Ethics of Belief Revisited", *American Philosophical Quarterly* 17, pp. 15-24.
- Mills, Eugene 1998: "The Unity of Justification", *Philosophy and Phenomenological Research* LVIII, pp. 27-50.
- Nozick, Robert 1993: *The Nature of Rationality* (Princeton: Princeton University Press).
- Papineau, David 1999: "Normativity and Judgement". *Proceedings of the Aristotelian Society*, supplementary vol.73, pp. 17-41.
- Parfit, Derek: *Rediscovering Reasons* (unpublished manuscript).
- Plantinga, Alvin 1993: *Warrant: The Current Debate*. (Oxford University Press, Oxford).
- Pollock, John 1987: *Contemporary Theories of Knowledge*. (Lanham, MD: Rowman and Littlefield).
- Putnam, Hilary 1983: "Why Reason Can't Be Naturalized", reprinted in his *Realism and Reason: Philosophical Papers*, vol. 3 (Cambridge: Cambridge University Press), pp. 229-247.

- Putnam, Hilary 1990: "Beyond the Fact/Value Dichotomy", in James Conant (ed.) *Realism with a Human Face* (Cambridge: Harvard University Press), pp. 135-141.
- Quine, W.V. 1986: "Reply to Morton White", in Hahn, L.E., and P.A. Schilpp (eds.), *The Philosophy of W.V. Quine*. (The Library of Living Philosophers, vol. XVIII) (La Salle, IL: Open Court) pp. 663-665.
- Quine, W.V. 1990: *Pursuit of Truth*. (Cambridge: Harvard University Press.)
- Railton, Peter 1997: "On the Hypothetical and Non-Hypothetical in Reasoning about Belief and Action", in Garrett Cullity and Berys Gaut (eds.) *Ethics and Practical Reason* (Oxford: Clarendon Press), pp. 53-80.
- Siegel, Harvey 1989: "Philosophy of Science Naturalized? Some Problems with Giere's Naturalism", *Studies in History and Philosophy of Science* 20, pp. 365-375.
- Siegel, Harvey 1990: "Laudan's Normative Naturalism", *Studies in History and Philosophy of Science* 21, pp. 295-313.
- Siegel, Harvey 1996: "Instrumental Rationality and Naturalized Philosophy of Science", *Philosophy of Science* 63, pp. 117-124.
- Sosa, Ernest 2002: "The Place of Truth in Epistemology". in Michael DePaul and Linda Zagzebski (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology* (Oxford: Oxford University Press).
- Velleman, David 2000a. "The Possibility of Practical Reason", reprinted in his *The Possibility of Practical Reason* (Oxford: Oxford University Press), pp. 170-199.
- Velleman, David 2000b. "On the Aim of Belief", reprinted in his *The Possibility of Practical Reason* (Oxford: Oxford University Press), pp. 244-281.
- Wedgwood, Ralph 2002: "The Aim of Belief", *Philosophical Perspectives* 16, pp. 267-297.
- Williams, Bernard 1973: "Deciding to Believe", reprinted in his *Problems of the Self*. (Cambridge: Cambridge University Press), pp. 136-151.
- Williams, Bernard 1981: "Internal and External Reasons", in his *Moral Luck* (Cambridge University Press, Cambridge), pp. 101-113.