

Chapter Fourteen

Is Truth the Primary Epistemic Goal?

Justification can take many forms: epistemic, prudential, moral. How can we distinguish between them? The standard answer is that epistemic justification is related to truth in a way in which prudential and moral justification are not. But how is epistemic justification related to truth? One answer says that epistemic justification is a means to achieving the goal of believing what is true and not believing what is false. In his defense of truth as the primary epistemic goal, Marian David examines a variety of goals: the truth goal, the knowledge goal, and the justification goal. He argues that the truth goal is more basic than the justification goal, but cautions against conceiving of the connection between truth and justification as a means–end relation. And, he suggests, if we reconceive the truth goal as the goal of having beliefs that are non-accidentally true, then we may view truth as the ultimate epistemic goal.

In his opposing essay, Jonathan Kvanvig defends the view that truth is not the primary epistemic value. In addition to truth, there are, according to Kvanvig, many other epistemic values: knowledge, understanding, wisdom, justification, and making sense of something. Among these, truth is not primary but just one goal among others. Truth becomes the primary goal only when we adopt a conception of epistemology as the analysis of knowledge. From Kvanvig's point of view, this is unduly reductive. Epistemology, Kvanvig suggests, is more than that: it is the study of successful cognition. Although he rejects the view that truth is the primary epistemic goal, in the end he endorses an account that is not so different from David's. The fundamental function of cognition, Kvanvig proposes, is that of determining, for any p , whether p . Nevertheless, he insists, there are epistemic goals which are not subservient to the truth goal. One such goal is that of epistemic responsibility, a goal that, according to Kvanvig, we can conceive of without invoking truth.

Truth Is Not the Primary Epistemic Goal

Jonathan L. Kvanvig

The question before us concerns the epistemic goal, standardly taken in epistemology over the past 50 years or so to be that of getting to the truth and avoiding error. In order to assess the plausibility of any answer to this question, it will be useful to begin by thinking about the question itself to make sure that it is properly understood. Asking about goals is asking about values or goods. To ask what one's goals are in, say, going to college, we inquire about the perceived values or goods that are being pursued. So the first thing to note about our question is that it asks about the values or goods that are epistemic in character.

The second point to note about the question is that it can be addressed from two quite different perspectives. One perspective is that of the theoretician. From this perspective, the question concerns what goods or values are central or primary for the theoretical task undertaken by the epistemologist, whatever that task may be. There is also another perspective, however, and that perspective is the point of view of those organisms about whose cognitive activity the epistemologist is theorizing. From this perspective, the question concerns the values or goods involved in the type of states and activities investigated by epistemologists.

It is important to keep the difference between these perspectives in mind when addressing this question, for there is no reason to assume that the answer to the question will be the same from either perspective. For example, it might be the case that truth is the primary good that defines the theoretical project of epistemology, yet it might also be the case that cognitive systems aim at a variety of values different from truth. Perhaps, for instance, they typically value well-being, or survival, or perhaps even reproductive success, with truth never playing much of a role at all.

Our question arises primarily from the perspective of the theoretical project of epistemology, and I will address the question from that point of view. We will see that the perspective of the cognitive system itself plays a role in this investigation, but only an ancillary one. In addressing the question before us I will be arguing for a negative answer to it. I will be arguing, instead, that there is a plurality of epistemic values and goals, and that though truth is an important epistemic goal, it has no claim to being the primary such value or goal. In order to argue for this plurality view, it is important to begin with a general account of the subject matter of epistemology, for the narrower one's conception of epistemology, the easier it is to defend the idea that truth is the primary epistemic goal. After explaining the appropriate domain of epistemological theorizing, I will take up the task of defending the pluralistic view.

What Is Epistemology?

Epistemology is often taken to be the theory of knowledge, but that conception is too narrow. At the most general level of characterization, epistemology is the study of certain aspects of our cognitive endeavors. In particular, it aims to investigate

successful cognition. Within its purview, then, are various kinds of cognizing, including processes such as thinking, inquiring, and reasoning; events such as changes in one's world view or the adoption of a different perspective on things; and states such as beliefs, assumptions, presuppositions, tenets, working hypotheses, and the like. Also within its purview is the variety of cognitive successes, including true beliefs and opinions, viewpoints that make sense of the course of experience, tenets that are empirically adequate, knowledge, understanding, theoretical wisdom, rational presuppositions, justified assumptions, working hypotheses likely to be true, responsible inquiry, and the like.

Two notes of caution are in order here. First, not just any kind of inquiry into kinds of cognizing and varieties of cognitive success counts as epistemology. Presumably, one can investigate these issues scientifically as well as philosophically, and I assume that they can be investigated in other ways as well (religiously, politically, morally, aesthetically, etc.). Epistemology is the result when these issues are investigated philosophically, but I will not here attempt any general characterization of the difference between philosophical and other forms of investigation.

The second point to note is that what kinds of success in cognition are relevant to epistemology is somewhat controversial. Beliefs that contribute to the well-being of the organism are successful in some sense of that term, and yet some will hold that such success is not the kind of success within the purview of the discipline of epistemology. Notice that success of a practical sort will advert to the causal consequences of holding the beliefs in question, and a common view is that epistemology is more concerned with intrinsic features of cognition, the kind reflected in talk of inquiry for its own sake. When we engage in inquiry for its own sake, successful results will partake of a kind of success that is independent of any causal contribution to well-being or other practical concerns. When epistemologists reflect on the nature of successful contribution and the extent to which an organism achieves it, the predominant approach has been to reflect on a kind of success that abstracts from the consequences of cognition, whether those consequences are practical, moral, religious, political, or social. I am inclined here to make a terminological restriction that what I mean by the use of "epistemology" and related terms is just this study of success which abstracts from the consequences of cognition, but I do not wish to be understood to denigrate more pragmatic approaches that claim to be epistemological yet deny that there is any value in, or value in thinking about, inquiry for its own sake. In part, my decision to limit a discussion of epistemology to one involving reflection on success of this rather rarefied sort rests on an attempt to consider the best case available to those who defend a positive answer to the question of whether truth is the primary epistemic goal. On these pragmatic approaches, such a positive answer has nothing to recommend it, but I will ignore such approaches here.

In slogan form, my characterization of epistemology is that it is the study of purely theoretical cognitive success, where the notion of what is purely theoretical is understood as above in terms of abstraction from the causal consequences of the success in question. As already noted, I grant that this characterization has a mild stipulative dimension in that it refuses to count as epistemology certain types of pragmatic approaches to the study of successful cognition. Even with this mild stipulation, however, the point I am trying to bring across is the breadth of this

conception of epistemology, breadth with ramifications for the question before us. Once we notice this breadth, we will be struck by the strong reductionist flavor of the claim that truth is the primary epistemic goal. One's first inclination should be to maintain that each independent kind of cognitive success within the purview of epistemology identifies a cognitive goal in its own right. From this viewpoint, epistemic goals include knowledge, understanding, wisdom, rationality, justification, sense-making, and empirically adequate theories in addition to getting to the truth and avoiding error. Once we have seen the variety of cognitive successes, the proper answer would seem to be that the class of epistemic goods is manifold, as wide as the class of cognitive successes. To answer otherwise would seem to engage in Procrusteanism involved in mild humor that begins with "there are two kinds of people in the world," except that this particular truncation would begin by trimming the number from two to one. Given this initial variety, how could the reductionist viewpoint be sustained?

Epistemic Values and Goals

One place to turn for a defense of the reductionist thesis is to the needs of epistemology itself. Perhaps in theorizing about cognitive success, there is a need for appeal to the goal of getting to the truth and avoiding error. I think something along these lines underlies the standard view that the epistemic goal is that of getting to the truth and avoiding error. The *Reader's Digest* condensed version of this approach goes like this: epistemology is the theory of knowledge, and knowledge includes, but is something more than, true belief. What more is required provides a connection between truth and belief. For example, it is well accepted that knowledge requires that the connection between belief and truth is *non-accidental*. Furthermore, for those who think that knowledge has a normative ingredient such as justification, rationality, responsibility, or the like, there will be a theoretical need to distinguish the sense of such terms from senses which have no place in an account of knowledge, such as practical senses. Perhaps that epistemic sense of the term should be characterized in terms of the goal, or good, of truth itself. If so, then perhaps the best way to think of the conditions needed for knowledge in addition to belief and truth are conditions best conceived in terms of some connection to truth.

Marian David provides an account of the matter along these lines. He says:

Knowledge, the epistemic concept par excellence, is usually defined in terms of belief, truth, and some other epistemic concept, say, justification: S knows p iff p is true, S belief p, and S is justified in believing p in a manner that meets a suitable anti-Gettier condition. Belief and truth, although fundamental to epistemology, are not themselves epistemic concepts. They are the nonepistemic ingredients in knowledge. This means that epistemology is not responsible for them; that is, as far as epistemology is concerned, belief and truth are *given* and can be invoked to account for epistemic concepts. The distinctly epistemic ingredient in knowledge is justification: the concept of S's *being justified* in believing p... Epistemology is certainly responsible for this concept. Indeed, once an account of knowledge is at hand, the task of epistemology pretty much reduces to the task of giving a theory of justification. (David, 2001)

David here endorses several of the claims noted earlier. First, his conclusion that epistemology's task reduces to that of giving a theory of justification requires the assumption that epistemology is the theory of knowledge. So David endorses

- (1) Epistemology is the theory of knowledge.

Second, David endorses the strong claim that

- (2) The only epistemic concept in an account of knowledge is justification.

David argues for (2) by claiming that belief and truth are not epistemic concepts, though he does not explain why they are not. Perhaps he is thinking as follows. Truth is a *semantic* notion and belief a *psychological* notion. As such, they should be investigated by semanticists and psychologists, not epistemologists.

Even if we accept this explanation, David's claim is a bit strong since it ignores the condition needed to solve the Gettier problem. (2) denies that attempts to solve the Gettier problem fall within the domain of epistemology, and that claim is mistaken. This point is more important than it might seem initially to be, since David's approach to showing that truth is the primary epistemic goal requires focusing exclusively on the theory of justification. Once we alter (2) to accommodate this point about the Gettier problem, no claims about the theory of justification can establish on their own the claim that truth is the primary epistemic goal.

I want to stress that we should not overestimate this difficulty, however. One approach to the Gettier problem, the defeasibility approach, explains the condition needed to handle that problem as a function on truth and justification. The simplest example of this approach says that having knowledge requires, in addition to justified true belief, there being no true information which, if learned, would result in the person in question no longer being justified in believing the claim in question. Though complications arise for this overly simple approach,¹ the point to note here is that the central concepts in this approach – truth and justification – already appear in the account of knowledge. If David's point can be sustained that justification is the only epistemic concept among the concepts of justification, truth, and belief, then this approach to the Gettier problem will be quite useful in the attempt to argue that truth is the primary epistemic goal by focusing on its role in the theory of justification. So even though David's account cannot succeed without some discussion of the Gettier problem, there is some hope for the idea that the Gettier problem will not force an alteration of his view that providing a theory of justification is the central task of epistemology. Perhaps David finds the defeasibility approach to the Gettier problem attractive, and it is this attraction that leads him to bypass any discussion of the Gettier problem in his defense of the claim that truth is the primary epistemic goal.

The defense of (2) is still troublesome even if we allow this explanation of why the Gettier problem does not undermine it. One problem with this explanation is that it proves too much. Justification is a property not only of beliefs but also of other things (such as actions) not within the domain of epistemology. According to the above explanation, if a concept is suitable for investigation by disciplines other than epistemology (such as semantics or psychology), then that concept is not an epistemic concept. Insofar as justification is also a property of actions, it is suitable for

investigation by action theorists, a subclass among philosophers of mind. So it appears that the above defense of (2) implies that there are *no* epistemic concepts in the proper account of the nature of knowledge!

One might reply that there is a distinctive concept of justification that is at home only in epistemology, and that concepts of justification which apply to other things such as actions are simply different concepts. In this way, the concept of justification needed in an account of knowledge is at home only in epistemology and is thus the only epistemic concept in an account of knowledge.

Such a view faces serious difficulties. Consider a simple juridical example such as the O.J. Simpson case. Some who watched the case closely say that contrary to the actual finding of not guilty, the evidence justified both their position that Simpson was guilty (beyond a reasonable doubt) and a decision by the judge to send him to prison for life. This claim predicates justification of two things: the first is a position of the speaker, that is, a belief, and the second is an action by a legal authority. On the view proposed in support of (2) above, we would need to treat this juridical pronouncement as semantically awkward, in the way it is semantically awkward to say that both sides of rivers and certain financial institutions are banks.

Let us put these difficulties aside for the moment, however, to see how (2) is supposed to take us to the idea that truth is the primary epistemic goal. On this issue, David provides us with two quite different answers. His first answer appeals to the idea of supervenience:

However, it is usually held that, at some point, the theory of justification has to “break out of the circle” of epistemic concepts and provide a nonepistemic “anchor” for justification by connecting it in some significant manner with nonepistemic concepts. ... It is not hard to see how the truth-goal fits into this picture. It promises to provide a connection between the concept of justification and the concept of true belief, tying together the different ingredients of knowledge. (David, 2001)

This attempt to find a central place for the truth goal in epistemology is unsuccessful. The need to anchor epistemic concepts in the non-epistemic is more plausibly taken as a demand that evaluative concepts supervene on non-evaluative concepts, and there is no reason to identify truth with the supervenience base of evaluative concepts in epistemology, any more than there is reason to identify the concept of the good as the supervenience base of evaluative concepts in ethics. For example, one answer to the demand for a supervenience base is broadly empiricist: justification supervenes on experiential states and logical (or quasi-logical) relations between propositions. In this view, what *makes* a belief justified is that it stands in the right logical or quasi-logical relationship to another belief or to (the propositional content of) an experience. One may be tempted to maintain that the concept of truth will still enter into this picture when we try to characterize logical relationships, but that reply is a red herring. First, logical relationships can be characterized syntactically as well as semantically. Second, and more important, even if truth is an ingredient in our account of this supervenience base, it is not present in the manner of a *goal* or anything of the sort. To argue that because truth is needed to explain logic, it must be that a primary epistemic goal would also commit one to the view that because the concept of a sensation is necessary to explain experience, sensation itself must be a primary epistemic goal according to empiricism.

There is a particular kind of approach to the nature of justification that fits David's picture a bit better than the empiricist example. Alvin Goldman's reliabilism of 1979 begins with the same demand to break out of the circle of epistemic concepts in order to clarify the nature of justification, and the way in which Goldman prefers to break out of the circle is in terms of the concept of reliability. The central non-epistemic concept in his theory is that of the reliability of a process or method of belief formation, the percentage of times that process or method generates a true belief. Given that it is percentage of truth over error by which we favor some processes and methods over others, we might wish to characterize our theoretical preference for some processes or methods over others in terms of the goal of truth over error, thereby securing the idea that truth is the primary epistemic goal.

This way of defending the claim that truth is the primary epistemic goal thus requires the subconclusion that reliabilism is the correct approach to the nature of justification, a position that is far from trouble free.² But the particular problems one might cite for reliabilism are not the central difficulty. The deeper problem is that the question we are attempting to answer arises at a different point in epistemological inquiry than the question of the adequacy of some particular theory of justification. The question whether truth is the primary epistemic goal is a meta-epistemological question, whereas the question of the adequacy of reliabilism is not. When we ask about the primary epistemic goal, we are asking whether there is a way of defending the idea that truth is the primary epistemic goal that is neutral between competing epistemological theories, and given that interpretation of the question, the approach that relies on the adequacy of a reliabilist approach in epistemology must be judged to be unsuccessful.

David seems to recognize the meta-epistemological character of the issue, for he provides another and quite different answer to the question of the relationship between justification and truth. He says:

Why is truth typically cast as a goal ...? Alston provides the reason. It is generally agreed that being justified is an evaluative concept of some sort: To say that believing *p* is justified or unjustified is to evaluate belief *p*, in some sense, as a good thing or as a bad thing, as having some positive status or some negative status. The suggestion is that this type of evaluation, epistemic evaluation, is most naturally understood along broadly teleological lines, as evaluating beliefs relative to the standard, or goal, of believing truth and avoiding error. (David, 2001)

David here endorses Alston's claim that the evaluative character of justification should be understood teleologically, and this explanation of the relationship between truth and justification is a suitably meta-epistemological account in contrast to the previous account which seems to depend on the particular epistemological theory of reliabilism.

What we need to know, however, is how this point about the teleological character of justification helps to show that truth is the primary epistemic goal. David's argument proceeds by contrasting the idea that truth is the goal with the idea that knowledge is the goal:

Although knowledge is certainly no less desirable than true belief, the knowledge-goal is at a disadvantage here because it does not fit into this picture in any helpful manner.

Invoking the knowledge-goal would insert the concept of knowledge right into the specification of the goal, which would then no longer provide an independent anchor for understanding epistemic concepts. In particular, any attempt to understand justification relative to the knowledge-goal would invert the explanatory direction and would make the whole approach circular and entirely unilluminating. After all, knowledge was supposed to be explained in terms of justification and not the other way round. This does not mean that it is wrong in general to talk of knowledge as a goal, nor does it mean that epistemologists do not desire to have knowledge. However, it does mean that it is bad epistemology to invoke the knowledge-goal as part of the theory of knowledge because it is quite useless for theoretical purposes: The knowledge-goal has no theoretical role to play *within* the theory of knowledge. (David, 2001)

David's argument thus proceeds as follows. We have already seen his claims that (1) epistemology is the theory of knowledge and (2) justification is the only epistemic concept within the theory of knowledge. These points imply that the central epistemic concepts are knowledge and justification, so the primary epistemic goal will have to be one related to these concepts. David's teleological conception of justification makes truth a contender for being the primary epistemic goal, and if the only epistemic concepts are justification and knowledge, the central competitor of the idea that truth is the primary epistemic goal would be that knowledge itself is the goal. Once we see the implication from (1) and (2) to the point that the central task of epistemology is to construct a theory of justification, the knowledge goal cannot be the primary epistemic goal on pain of rendering our epistemology viciously circular. Only the truth goal can be central if our theory has any hope of adequacy, so truth must be the primary epistemic goal.

The form of argument here bears scrutiny since the particular argument in question relies on the false assumption that epistemology is the theory of knowledge. Given this assumption, the competitor to the idea that truth is the primary epistemic goal is the idea that knowledge is also an epistemic goal. Given the more general conception of epistemology outlined earlier and the variety of cognitive successes it implies, we will need a generalization of David's argument in order to conclude that truth is the primary epistemic goal. Such a generalization could begin by allowing that any purely theoretical cognitive success is of value and hence a suitable epistemic goal. But for each such goal, it has no theoretical goal to play within the project of theorizing about it, for such an account would make the theory "circular and entirely unilluminating." Moreover, it would have to be claimed, any goal that would render a theory circular and unilluminating in this way cannot be an epistemic goal, leaving truth as the only standing candidate for that role.

So, on this view, nothing of epistemic value can be an epistemic goal. Instead, only values that fall outside the domain of cognitive successes can legitimately be cited when explaining epistemic success. It also follows from this point that truth and belief cannot themselves be epistemically valuable, for otherwise getting to the truth cannot itself be an epistemic goal.

Contrast this result with an alternative picture. For simplicity, let us grant David's claims that epistemology is the theory of knowledge and that knowledge is to be understood in terms of justified true belief plus a Gettier condition. The alternative picture holds that each element in the theory of knowledge should isolate some epistemic good (something whose value contributes to and helps explain the value of knowledge). So, for example, truth itself must be an epistemic good, implying that it

has value exceeding that of mere empirical adequacy. In a similar vein, belief must be an epistemic good, implying that Pyrrhonian counsel to abandon beliefs in favor of acquiescing to the appearances cannot, in any literal sense, accord with the facts about what is important from a purely theoretical point of view. Similar points could be voiced about the other two conditions as well, and the accumulated result of these points is that if truth is an epistemic goal so are the other values just mentioned. As a result, truth is not the primary epistemic goal, but rather one among several epistemic values that have equal claim to being epistemic goals. This alternative picture strikes me as much more plausible than the idea that nothing that is epistemically valuable can itself be an epistemic goal or good.

The generalization of David's argument above has a further problem. This generalization requires that for any epistemic value *V* other than truth itself, the fundamental goal in terms of which a theory of *V* is constructed is truth itself. Once we put the view in these terms, I think it is fairly easy to see that it is false, though it may be more plausible in some parts of epistemology than in others. Perhaps, for example, it has some plausibility with respect to the value of belief itself. One might maintain that the value of belief is in some way dependent on the value of truth in that cognitive activity in general is prompted by the *desire* to get to the truth, and that belief is the result of such alethic motivation. Alternative alethic accounts are also available, ones which posit something sub-intentional as the driving force behind such activity, perhaps a need or interest or drive, or perhaps even an instinct, to get to the truth. One might even hold that it is the *function* of cognition to get to the truth and avoid error, whether the organism in question can be properly characterized by any of the intentional or sub-intentional descriptions just noted.

Such a view is not without difficulties, however, for there are alternatives to this truth-based account with some plausibility. To the extent that an organism has explicitly intentional motivation for cognitive activity, it may be more plausible to characterize those motivations in terms of a desire for knowledge or understanding rather than in terms of a desire to get to the truth, for we typically view mere true belief as something less than what we really want. Moreover, this point about fully intentional cognitive activity carries over to subintentional contexts as well. If the function of cognition is to be given a factive characterization – one which implies truth, as opposed to one which requires only adequate coping with an environment for purposes of survival and reproduction – the factives of knowledge and understanding have at least as much going for them as truth itself. That is, it may be more plausible to hold that the function of cognition is to produce knowledge or understanding of one's surroundings than it is to maintain that the function of cognition is to get to the truth about one's surroundings.

I do not wish to rest my negative answer to the question of whether truth is the primary epistemic goal on a rejection of truth-based accounts of cognition, however. My own view is that something close to a truth-based account is correct. Since my answer to the question before us does not depend on my view of this matter, I will present only the briefest account of it here. In my view, the fundamental function of cognition and the fundamental intentional attitudes involved in cognitive activity can be characterized without appeal to the concept of truth. The function in question is that of determining, for any claim *p*, whether *p*: for example, whether it is raining, or whether a particular object is a predator. Such a characterization appeals to no factive concepts stronger than truth, however. Moreover, even though the affective states

underlying cognitive activity often involve factives such as knowledge or understanding, such states are not universal – not everyone wants knowledge, for example, and not everyone is motivated by a concern for understanding. What is fundamental to explicitly intentional cognitive activity as such is a weaker state, for even small children and animals engage in fully intentional cognitive activity. We characterize curiosity as the desire to know, but small children, lacking the concept of knowledge, display curiosity nonetheless. The conclusion to draw from these considerations is that an adequate characterization of cognition as such will appeal to no motivational structure, either intentional or non-intentional, involving a factive element stronger than truth. To the extent, then, that we are persuaded that such a motivational structure must involve facticity of some sort, we should have no objection to characterizing cognition in terms of the goal of truth.³

Let us return to the main issue, arising from the fact that the argument we are considering requires that for any epistemic value *V* other than truth itself, the fundamental goal in terms of which a theory of *V* is constructed is truth itself.

Even if we grant that truth is in some way fundamental to our understanding of cognition itself, much more is needed to sustain this generalization. In particular, one would have to be able to argue that truth is fundamental to any of the kinds of purely theoretical cognitive successes noted earlier. Some such successes clearly involve the concept of truth, concepts such as knowledge, understanding, wisdom, likely hypotheses, and justified beliefs; and if these concepts involve the concept of truth, there is some hope that the goal of truth is somehow fundamental to an adequate account of these concepts. There are also cognitive successes that are not obviously truth related, such as the concepts of making sense of the course of experience and having found an empirically adequate theory. Both of these concepts can be explained without recourse to the goal of truth. An empirically adequate theory is one that will never be refuted by the course of experience, and one makes sense of the course of experience by developing a classification system for experiences together with a theory of explanation of how the various categories are explanatorily related.

Another example of a cognitive success that is not obviously truth related is that of responsible inquiry. Explaining why will take a bit of work, and we can begin by distinguishing deontic concepts from evaluative ones. Most ethical theories are teleological, explaining deontic notions such as obligation, permission, and forbiddenness in terms of evaluative notions such as the good. Not all do, however. Deontological theories refuse such definitions, explaining deontic concepts in other ways, the paradigm example being a Kantian theory that explains forbiddenness in terms of internal contradictions between universal maxims.

Epistemology has seen defenders of what is termed deontologism, and here Chisholm is the paradigm example. Epistemic deontologists have not been deontologists in the sense defined in the previous paragraph, however, for they have defined epistemic duties in terms of epistemic values such as truth. Chisholm, for example, says our epistemic duties arise from the fundamental duty to (do one's best to) believe the truth and avoid error (Chisholm, 1977).

What would real deontologism look like in epistemology? A good example of such can be developed along Bayesian lines. According to one kind of Bayesianism,⁴ each of us has a complete theory of evidence in our heads, encoded in the form of conditional

probabilities. Such a theory can explain the forbiddenness of a (degree of) belief in terms of logical notions such as logical inconsistency, rather than in terms of value terms. A degree of belief is forbidden when, in the face of new experience, it becomes logically inconsistent with the relevant conditional probabilities.

Other examples are available as well. Though Richard Foley⁵ wholeheartedly endorses the standard epistemic goal of getting to the truth (now) and avoiding error (now), his theory doesn't require this teleological aspect. On his theory, a belief is justified if and only if it conforms to one's deepest epistemic standards. These standards are epistemic principles one would endorse given as much time to reflect as is needed to reach a stable point of view.⁶ Once we have such internal standards playing a theoretical role, real deontology can emerge, for it is in conforming to these standards that one's epistemic duties are satisfied.

Teleology in epistemology is needed only when the epistemic principles are licensed in some other, non-subjective way. If one's theory is truly subjective, then consistency with an internalized theory of evidence is all that is needed for defining key epistemic concepts such as justification. If, however, the theory of evidence is objective, then one needs some standard against which to assess its adequacy. It is here that epistemic teleology has its natural home.

Given this background, let us return again to the concept of responsible inquiry. The concept of epistemic responsibility can have strong subjective overtones, though it need not. A theory of epistemic responsibility may be modeled on the ethical notion of responsibility, which is to be contrasted with the notions of praiseworthiness and blameworthiness. In the moral sphere, one can be blameless in failing to live up to one's responsibility, and if we analogize to epistemic responsibility, we will be willing to say that one's inquiry might be irresponsible but blameless. If we refuse to distinguish epistemic responsibility from epistemic blamelessness, our theory of epistemic responsibility will be much more subjective, and the possibility of a non-teleological account of it will emerge.

Not much turns on this issue in the present context, however, for even in the face of a compelling argument that responsibility should be identified with blamelessness, there will still be an epistemic concept of value that can be understood in non-teleological terms – namely, intellectually blameless inquiry. So not only are there the notions of sense-making and empirical adequacy that do not presume the goal of truth, there are other, normative notions that can equally be clarified without appeal to the goal of truth.

Conclusion

The slogan that truth is the epistemic goal has influenced much of epistemology in recent history, and we have seen that there is some truth in this idea. In particular, the idea of truth as a goal will play a significant and fundamental role in our understanding of cognition itself. The claim that truth is the primary epistemic goal, however, goes well beyond such a claim. It implies that the concept of truth plays a fundamental role in any adequate theory of any cognitive success of a purely theoretical sort. The view I have defended is that such a position has too narrow a view of the variety of epistemic values and goods, and that once we appreciate this variety and the broadened

conception of the domain of epistemological inquiry, we will go no further than to embrace the idea that truth is an important and central epistemic value and goal.

Notes

- 1 For discussion of these complications, see Klein (1981).
- 2 To my mind, the most serious such problem is that reliabilism cannot explain adequately the concept of propositional justification, the kind of justification one might have for a proposition one does not believe or which one disbelieves (believes the opposite). I have argued this point in several places (e.g., Kvanvig, 1990, 1992, 1996, 2000, 2003).
- 3 For a more complete investigation of these issues, see the chapter on the value of truth in my 2003.
- 4 For an intuitive account of Bayesianism and its commitments, see van Fraassen (1990).
- 5 The view I discuss here is from Foley (1986).
- 6 This “operational” definition of the concept of one’s deep standards leads to various objections to Foley’s view, of the sort that usually plague counterfactual accounts. For example, if one’s deepest standards are chary of reflection of any kind, what results from Foley-reflection will be self-stultifying. To avoid these problems, one could develop a Foley-like view where the counterfactuals in question were taken to be evidentially related to the question of what one’s deep standards are, but deny that they are related definitionally.

References

- Chisholm, R. (1977) *Theory of Knowledge*, 2nd edn. Englewood Cliffs, NJ: Prentice Hall.
- David, M. (2001) Truth as the epistemic goal. In *Knowledge, Truth, and Duty. Essays on Epistemic Justification, Virtue, and Responsibility*, ed. M. Steup. Oxford: Oxford University Press.
- Foley, R. (1986) *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.
- Goldman, A. (1979) What is justified belief? In *Justification and Knowledge*, ed. G. Pappas. Dordrecht: Reidel.
- Klein, P. (1981) *Certainty*. Minneapolis: University of Minnesota Press.
- Kvanvig, J. (1990) The basic notion of justification. Christopher Menzel, co-author. *Philosophical Studies* 59: 235–261.
- Kvanvig, J. (1992) *The Intellectual Virtues and the Life of the Mind: On the Place of the Virtues in Contemporary Epistemology*. Savage, MD: Rowman & Littlefield.
- Kvanvig, J. (1996) Plantinga’s proper function theory of warrant. In *Warrant and Contemporary Epistemology*, ed. J. Kvanvig. Savage, MD: Rowman & Littlefield.
- Kvanvig, J. (2000) Zagzebski on justification. *Philosophy and Phenomenological Research* 60: 191–196.
- Kvanvig, J. (2003) Propositionalism and the perspectival character of justification. *American Philosophical Quarterly* 40 (1): 3–18.
- Kvanvig, J. (2007) *The Value of Knowledge and the Pursuit of Understanding*. New York: Cambridge University Press.
- Pappas, G. (ed.) (1979) *Justification and Knowledge*. Dordrecht: Reidel.
- Steup, M. (ed.) (2001) *Knowledge, Truth, and Duty. Essays on Epistemic Justification, Virtue, and Responsibility*. Oxford: Oxford University Press.
- van Fraassen, B. (1990) *Laws and Symmetry*. Oxford: Oxford University Press.

Truth as the Primary Epistemic Goal: A Working Hypothesis

Marian David

I Possessing Truth

Let us first consider what it means to talk of truth as a goal, without worrying about its being the primary epistemic goal. "I want truth" is a bit like "I want fruit." Like fruit, truth comes in pieces. We have separate words for the different sorts of pieces fruit comes in. The pieces truth comes in we can just call *truths*. Philosophers often call them true *propositions*. The term is convenient because it also fits the pieces falsehood comes in. Say you are looking at the monitor and assert that the flight is leaving. I am looking too and I also assert that the flight is leaving. We both assert the same thing, the same proposition: the proposition that the flight is leaving. The proposition is either true or false. If it is true, we are both right; if it is false, we are both wrong. So, "I want truth" says that I want true propositions. "I want *the* truth," taken literally, says that I want exactly the truth, the whole truth and nothing but the truth; that is, taken literally, it says that I want all the true propositions and only the true propositions, no false ones. Of course, normally one would not take it literally; one would assume that what I really wanted is all and only the true propositions relevant to some salient subject matter not explicitly specified.

"I want truth" and "I want the truth" do not mention what I want with the truth. But it is clear that I want at least to have it, to possess it. How does one possess truth? We can possess truth by way of *believing* it, by believing true propositions. We can possess truth by way of *knowing* it, by knowing true propositions. "I want truth" can be understood either way. So the broad idea of truth being a goal comes in two versions, corresponding to the two ways of possessing truth. Let us call them the *true-belief goal* and the *knowledge goal* respectively. They are both truth goals.

On the believing-way of possessing truth, "I want truth" comes out as: "I want to believe true propositions." The more colloquial alternative, "I want to have true beliefs," is fine too as long as we remember that using the word "belief" for the things possessed can be awkward if believing is the mode of possessing them. We would not want to say: "I want to believe true beliefs." The noun "belief" is, as one puts it, act/object ambiguous. It might be used to refer to the act, or attitude, of believing a proposition. It might also be used to refer to the object, or content, of the attitude, to the proposition believed. If you believe that the flight is leaving, and someone says that your belief is true, then the noun "belief" refers to the proposition you believe. If someone were to say that your belief was silly, the noun would refer to your attitude of believing the proposition. The desire for truth relevant to the true-belief goal is a desire directed more at the attitude than the propositions; it is a desire for believing what is true, rather than for the truth of what one believes.

Propositional knowledge entails truth. One cannot know a proposition unless it is true: you cannot know that the flight is leaving at eight, unless it is leaving at eight

(of course, you might *think* you know it is leaving at eight even though it is not). Since propositional knowledge entails truth, we can express the knowledge goal simply as “the goal of having knowledge,” without mentioning truth. It is a truth goal nevertheless: it is the goal of possessing truths by way of knowing them.

The true-belief goal and the knowledge goal are both truth goals. Are there really two distinct goals here? Propositional knowledge entails true belief: one cannot know a proposition unless one believes it. (Sure enough, we sometimes say things like: “I don’t believe it; I know it.” But this is short for: “I don’t *merely* believe it; I know it.” We wouldn’t say: “I know it, but I don’t believe it.”) Since propositional knowledge entails true belief, the goal of having knowledge and the goal of having true beliefs are not entirely distinct. They are at least “overlapping” goals in this sense: if you want to have knowledge, you want something you cannot have without having true beliefs.

Epistemologists like to emphasize that knowledge requires something in addition to true belief: justified or warranted or rational belief, or belief based on good reasons or on adequate evidence (and they usually add that more is required still). This seems right for many typical uses of the word “know.” There also seem to be uses of “know” that do not require more than true belief. Alvin Goldman (1999) calls this the *weak* sense of “knowledge”: when an unfaithful husband greets his lover with the words “She knows,” he is not worried about whether his wife is justified in believing that he is cheating on her; all he is worried about is that she believes it and that it is true. The goal of having knowledge in this weak sense coincides with the goal of having true beliefs. But let us reserve the term “knowledge” for the *strong* sense, the sense in which it requires more than true belief.

We can say, then, that the knowledge goal and the true-belief goal are different, albeit overlapping, truth goals.

Invocations of truth as a goal, as something desired or aimed at, show up fairly frequently, especially in religious, philosophical, and broadly scientific contexts: phrases like “the search for truth” and “the pursuit of truth” are almost commonplaces there. Such invocations usually leave open which of the two truth goals is intended. In general, any reference to truth as a goal might be taken to refer to the goal of believing truths or to the goal of knowing truths – our title is a case in point – and even when the word “knowledge” is used, it is often difficult to tell whether it is intended in the weak or the strong sense. Compare William James’s (1911) commandment, “We must know the truth; and we must avoid error”; which he rephrases almost immediately as “Believe truth! Shun error!”

II Truth Goals

The goal of having true beliefs, the goal of believing true propositions, is an indefinite or indeterminate goal: it is vague about how much of the truth is being aimed at. The most ambitious determinate goal would be the one corresponding to “I want to believe *the* truth”; it would be the goal of believing all the true propositions and not believing any false ones. Using “ $\forall p$ ” to abbreviate the universal quantifier “for all propositions p ,” it might be represented like this:

$$(1) \quad G (\forall p) (Tp \rightarrow Bp \ \& \ Bp \rightarrow Tp),$$

but we must keep in mind that the relevant desire is supposed to be directed at the attitude rather than the propositions. So the goal represented in (1) should be understood in terms of wanting to believe all and only true propositions, or wanting to be such that one believes all and only true propositions, rather than wanting all propositions to be such that they are true if and only if one believes them.

It is fairly clear that we don't actually have this goal. The first part, the part about believing all the truths, looks like an absurd thing to want – partly because it is too obvious that there must be way too many truths for it to be humanly possible to believe them all, partly because there are vast numbers of unimportant and boring propositions, which, it seems, we wouldn't particularly want to believe even if they were true (and the truths are so redundant: any conjunction of true propositions is itself a true proposition). The second part, though, believing only truths, not believing any falsehoods, does look like something we want – though we might say "I would like that," rather than "I want that," signaling that we don't think the chances of reaching this state are at all good. Still, because of its first part, (1) looks like an overly ambitious goal – verbal "commitment" to it ("Tell me the truth!") would normally be neither taken nor meant literally. One would automatically assume that it was meant to be restricted to all the relevant truths about some salient subject matter not explicitly specified. (But we should also remember that theists typically attribute omniscience to God, which involves believing all and only the truths. This indicates that we do tend to regard this state as something valuable, as some sort of ideal, albeit a remote one.)

Explicitly restricting the first part of (1) to "important and interesting" propositions would yield a goal that might be ascribed to us with some plausibility.

The indefinite goal mentioned earlier, the goal of "having true beliefs," can be regarded as a tacitly as well as indeterminately restricted version of (1). Indeed, this indefinite goal may well be the one most plausibly ascribed to us; it looks like a goal we actually have – maybe just because it is so vague. It is naturally understood with a suppressed rider about falsehood. Say like this: it is the goal of having true beliefs and few, or no, false ones (or maybe the goal of having a large stock of true beliefs containing few, or no, false ones; or maybe simply the goal of having true beliefs rather than false ones).

There might be some propositions that are so discomfoting (painful, upsetting, disgusting) to believe that we wouldn't want to believe them even if they be true, and some so comforting to believe that we would want to believe them even if they be false. Would this show that we don't have the indefinite goal? No. It would not even show that we don't have unrestricted (1) as a goal. We have various goals. In some cases different goals come into conflict and one "loses out." In some cases a goal can be "neutralized" by contrary emotions. This does not show that we don't have the beaten or neutralized goal; it merely shows that the goal is not an absolute one.

The goals considered so far are "collective" goals. They refer not to individual propositions but to whole collections of propositions: for example, unrestricted (1) refers to the collection containing all and only the true propositions; the indefinite goal refers to some "fuzzy" collection, or fuzzily refers to a collection. But talk about wanting to have true beliefs can also be intended "distributively." Consider the generalization:

$$(2) (\forall p) G (Tp \rightarrow Bp \ \& \ Bp \rightarrow Tp),$$

which says that, for every proposition p , one has the goal (wants) to believe p if and only if p is true. Whereas (1) represents a single goal referring to a collection of propositions, (2) itself does not actually purport to represent any goal at all; rather, it represents a (false) statement ascribing to us a huge number of individual relativized goals, one for each proposition p . It is clear that we can't have all these goals: there are way too many propositions.

But we can at least say this. For each proposition p such that we are seriously asking whether p is true, we will have the relativized true-belief goal with respect to p ; that is, for each such proposition p , we have the goal (want) to believe p if and only if p is true: the very fact that we are seriously asking (ourselves or others) whether p is true shows that this is one of our goals with respect to p , at least at the time at which we are asking. Actually, what we want with respect to any such p is more completely captured like this: to believe p , if and only if p is true, and to believe $\neg p$, if and only if $\neg p$ is true, where " $\neg p$ " refers to the negation of p . (Unlike Williams (1978, p. 38) and Sosa (2003), I take the "only if" parts to be important. Assume that, having asked whether p is true, you are being told that p is true and also, maybe by someone else, that $\neg p$ is true. Since one of them must be true, you are being told the truth with respect to $p/\neg p$. But this is not what you wanted. You wanted *only* the truth with respect to $p/\neg p$: you did not want to get both, the truth and the falsehood.)

Note that there is something funny about how collective true-belief goals relate to goals concerning individual propositions. Say you are looking through a list of bargain CDs, wanting to buy all but only the ones by Eminem. You have reached the next item; it is called *Slim Shady*. You will want to buy *Slim Shady* if and only if it is by Eminem. Now assume you come to believe that *Slim Shady* is by Eminem: then you will (probably) want to buy *Slim Shady* (unless you believe that getting it would interfere with a more important goal, like not spending more than a certain amount). Here the collective goal relates fairly straightforwardly to a non-conditional goal with respect to one particular thing via your belief that the thing belongs to the collection of things you want. Now assume you are going through a list of propositions, wanting to believe all and only the ones that are true. You have reached the next item: the proposition p . Like before, we can say that you will want to believe p if and only if it is true. But now assume you come to believe that p is true. Can we again go on and say, non-conditionally, that you will (probably) want to believe p ? This seems odd at best. Having come to believe that p is true, you already believe p (it is hard to see how you could believe that p is true without believing p); and once you believe p , it is at best odd to say that you *want* to believe p . Sosa (2003) suggests this shows that collective true-belief goals can play no role in "guiding" us about what to believe – whereas the collective Eminem goal can play a role in guiding you about what record to buy.

Let us briefly consider the knowledge goal too. Begin with the ambitious goal of having all knowledge; that is, the goal of knowing everything. To know everything would be to know everything there is to know. What is there to know? Truths, that is, true propositions, but no falsehoods: a false proposition cannot possibly be known (although it can be known to be false, but that would be knowing a truth about a falsehood). So part of the ambitious goal can be represented as " $(\forall p)(Tp \rightarrow Kp)$," which brings out that the knowledge goal is directed at possessing truths, just like the true-belief goal. What about a second, negative, part for the knowledge goal, a part corresponding to the $(Bp \rightarrow Tp)$ -part of (1)? I am not quite sure. There seem to be two options:

- (3.1) $G (\forall p) (Tp \rightarrow Kp \ \& \ Bp \rightarrow Tp)$,
 (3.2) $G (\forall p) (Tp \rightarrow Kp \ \& \ Bp \rightarrow Kp)$.

According to (3.1), the ambitious knowledge goal would be the goal of knowing all the truths and not believing any falsehoods (avoiding error). But one might hold that the second part of the goal should be stronger than the second part of (1); that it should be about not believing something unless one knows it. Accordingly, (3.2) describes the goal in terms of wanting to be such that one knows all the truths and believes only what one knows. Note that, just as the condition $(Tp \rightarrow Kp)$ subsumes the condition $(Tp \rightarrow Bp)$, due to knowledge requiring true belief, so the condition $(Bp \rightarrow Kp)$ subsumes the condition $(Bp \rightarrow Tp)$, and for the same reason.

Problems about the plausibility, or rather implausibility, of thinking that the (3)s describe goals we actually have arise here in just the same way as they did with respect to unrestricted (1). To get goals more plausibly ascribed to us, the first parts of the (3)s would have to be restricted to important and interesting propositions. The collective knowledge goal most plausibly ascribed to us may again be an indefinite one, say, the goal of having knowledge (and avoiding error, or not believing things unless one knows them); which can again be regarded as some sort of tacit fuzzy weakening of an ambitious goal. And again collective goals should be distinguished from relativized goals:

- (4.1) $(\forall p) G (Tp \rightarrow Kp \ \& \ Bp \rightarrow Tp)$,
 (4.2) $(\forall p) G (Tp \rightarrow Kp \ \& \ Bp \rightarrow Kp)$.

These statements ascribe goals that are relativized to individual propositions p . Of course, no one can have all these goals, one for each proposition p . But we can say that, for each proposition p such that we want to know whether p is true, we will have such a relativized goal with respect to p – again, the full goals we have with respect to such p 's should add clauses about wanting to know $\neg p$, if $\neg p$ is true.

III Epistemic Goals

Among the many goals we have, or might have, only some appear to be directly relevant to the subject matter of epistemology. Let us call them *epistemic* goals. The primary epistemic goal would then be the goal that is the most important or most fundamental of the goals directly relevant to the subject matter of epistemology. What are the candidates? Epistemology is the theory of knowledge. So the main candidates should be the goal of having knowledge, the goal of having true beliefs, and some goal relating to the other main requirement for knowledge, say *having justified beliefs* (or warranted or rational beliefs, or beliefs based on good reasons or on adequate evidence – for convenience, I will often use the concept of justified belief as representative of this group).

It is difficult to be precise about the goal of having justified beliefs – in part that is because there is little agreement about what justification actually amounts to. Maybe the following will suffice for the indefinite version of the goal: it is the goal of believing a proposition, if one has justification for believing it, and only if one has justification

for believing it (but we should not assume that one *has* a justification only if one has *given* one's justification). Alternatively, we might say that it is the goal of believing a proposition, if one has adequate evidence for it, and only if one has adequate evidence for it. Again we should distinguish between the single collective goal of having justified beliefs and the relativized goals we have, or might have, with respect to individual propositions *p*, to believe *p* if and only if we have justification for believing *p*.

The goal of having knowledge and the goal of having true beliefs are both truth goals. What about the goal of having justified beliefs? Is it also a truth goal? No: having justified belief is not a way of possessing truth. This is because, unlike true belief and knowledge, justification does not entail truth: a belief can be justified even though it is false. This does not mean that there aren't some forms of justification that do entail truth – justification by mathematical proof might be one. Such forms of justification are called “infallible” justification. But most forms of justification are fallible: they don't entail (guarantee) truth, e.g. justification by induction, or by the evidence of our senses. So, in general, having a justified belief does not entail that the justified belief is true.

We have identified three main epistemic goals: having justified beliefs, having true beliefs, and having knowledge. Which one, if any, can be regarded as the primary epistemic goal?

Many of our goals are connected. Often we want something because we want something else and because we realize that if we get the former we will (probably) get the latter, or because we realize that without getting the former we can't get the latter. Chains of goals we have because of other goals we have terminate in “basic” or “intrinsic goals,” things we want, or value, just for their own sake and not for the sake of anything else. So, considering our epistemic goals, we might ask whether our epistemic goals are connected in some such way. If there is a single one at the bottom of such a chain, then that would be the primary epistemic goal.

Note that to claim that something is the primary *epistemic* goal does not imply that it is a basic goal all things considered, that we want it only for its own sake, that it is an “absolutely” intrinsic goal. Take the idea that possessing truths (by believing them? by knowing them?) is the primary epistemic goal. We want to possess truths in part for their own sake (we tend to be simply curious about things) and in part, or maybe mostly, because we think that having truths will increase our chances for securing our other, non-epistemic goals (in particular, we want truths about the most effective means for securing our other goals). The idea that possessing truths is the primary epistemic goal does not commit itself on this issue beyond the claim that there are no other *epistemic* goals for the sake of which we want truths – never mind any non-epistemic goals for the sake of which we might want truths. It only commits itself to the claim that possessing truths is a “relatively” intrinsic goal, that it behaves like an intrinsic goal within the domain of epistemology.

Actually, it commits itself to even less, or rather, to something slightly different. The idea that something is the primary epistemic goal has more the force of an “ought” than the force of an “is.” It says that, to the extent that we do have epistemic goals, we *ought* to have them because we have the one singled out as primary. Obviously, it is rather risky to make claims about what we actually want and about why we want what we want, without asking or studying many of “us.” Even if all of us do have the epistemic goals discussed earlier (or at least closely related ones), and even if these goals

are in fact connected in the manner mentioned above, some of us might not realize that their epistemic goals are so connected. It would then be wrong to say that *we* have the other goals because we have the goal that is in fact the primary one; all that can be said is that we ought to. One can then go on and say that, to the extent that we realize this, we will tend to have the other epistemic goals because we have the primary one.

IV The True-belief Goal and the Justified-belief Goal

The goal of having justified beliefs is not a truth goal: having justified beliefs is not a way of possessing truths. Still, it is very plausible to think that it is a truth-oriented goal in the sense that it derives from, or depends on, a truth goal. The truth goal most often mentioned in this connection is the goal of having true beliefs. Consider the following claims:

- A1 If you want to have TBs you ought to have JB's.
- A2 We want to have JB's because we want to have TB's.
- B1 If you want to have JB's you ought to have TB's.
- B2 We want to have TB's because we want to have JB's.

Assume we want to have true beliefs and want to have justified beliefs. A1 would give us a reason for connecting these goals in the manner described in A2: if we believe what A1 says, we ought to and probably will connect these goals in the manner described in A2. If, on the other hand, we believed what B1 says, we ought to and probably would connect these goals in the manner described in B2.

Comparing these pairs, it is obvious that the A's are way more plausible than the B's. Indeed, initially one may even think that the B's have nothing going for them at all, that they are just false. That is not quite right, however. True beliefs are generally useful; hence, they can be instruments for acquiring new justified beliefs as well as for acquiring justifications for old beliefs, provided they are true beliefs of the right sort: namely, true beliefs about how evidence, old or newly acquired, bears on propositions we don't yet believe; true beliefs about how the evidence we already possess bears on our old beliefs; true beliefs about where to find new evidence for or against old beliefs; and so on. (If you come to believe, correctly, that some evidence *e* you have just acquired adequately supports some proposition *p*, and you then come to believe *p* in part because of your true belief about how *e* bears on *p*, then this belief contributed causally to the fact that you now hold the justified belief *p* – this is not the typical case though: more typically, acquiring *e* will just make you believe *p*, without going through the higher-level belief about *e*'s bearing on *p*). True beliefs of this special sort – roughly, true beliefs about epistemic matters – have a certain measure of instrumental value relative to the JB goal. They have “epistemic utility” (Firth, 1998b) insofar as they tend to have causal consequences (namely the acquisition of justified beliefs) that contribute to the realization of one of our epistemic goals. So, the Bs are not simply wrong. Still, compared to the As, they look tenuous. They apply only to true beliefs of a special sort, whereas the As apply quite generally to justified beliefs of all sorts.

Moreover, the rather weak reason B1 provides for wanting to have true beliefs because we want to have justified beliefs seems itself to take us back to the desire for true beliefs: to some extent we want true beliefs (namely true beliefs about epistemic matters) because we want justified beliefs (about many things); and we want justified beliefs (about many things) because we want to have true beliefs (about many things). In short: we want true beliefs about epistemic matters because we want true beliefs about many things. The As have much more going for them than the Bs.

The claims above are put in terms of collective epistemic goals. The B's are not entirely wrong because, collectively speaking, having true beliefs tends to have some epistemic utility relative to the JB goal: to the extent that a collection of true beliefs contains special true beliefs about epistemic matters, the collection will tend to promote the acquisition of justified beliefs. The As, on the other hand, are much more plausible than the Bs, because they stay plausible even when taken as claims about particular goals relativized to individual propositions p – read “iff” as “if and only if”:

- A3 For any p : if you want to <believe p iff p is true>, then you ought to <believe p iff you have justification for believing p >.

The corresponding claim for B1, the claim that results from exchanging the parentheses, has no plausibility at all. Note that A3 does not talk about the instrumental value of having justified belief in p relative to the goal of having true beliefs. It does not talk about the causal consequences justified belief in p has, or tends to have, for acquiring other beliefs that are true: it does not talk about any causal consequences of having justified belief in p at all. Where, then, does A3 “come from,” as it were, if not from the epistemic utility of having justified beliefs vis-à-vis the goal of getting true beliefs? The intuitively most plausible answer to this question is surely something like:

- A4 You ought to have justified beliefs rather than unjustified beliefs, given that you want to have true beliefs rather than false ones, because justified beliefs are likely to be true, at least considerably more likely to be true than unjustified ones.

The intuitive plausibility of A1 is grounded largely in A3, which seems to rest on A4. Note again that A4 does not talk about the causal consequences of having justified beliefs. It says that justified beliefs *themselves* are likely to be true. This claim is much stronger than the claim that, collectively speaking, having justified beliefs tends to make it likely that one has true beliefs. The latter would hold even if having justified beliefs merely tended to promote the subsequent acquisition of true beliefs. Although, somewhat confusingly, having justified beliefs does tend to do that too. In addition to A4 (justified beliefs being themselves likely to be true), there is also an “epistemic utility”-aspect to having justified beliefs, but this aspect depends on A4. Justification is transferable: to the extent that you acquire new beliefs properly based on justified beliefs, your new beliefs will themselves be justified; hence, because of A4, they will themselves be likely to be true. So, because of A4, having justified beliefs tends to have a measure of general epistemic utility relative to the TB goal: it tends to promote the growth of true beliefs. This epistemic utility of having justified beliefs, which itself depends in part on A4, adds to the plausibility of A1.

It is tempting to talk of justification as if it were essentially a *means* to the truth (e.g., BonJour, 1985, pp. 7–8). But note that such talk can be seriously misleading. It connects with the *less* important aspect, the “epistemic utility”-aspect, underlying A1. Though it is correct to say that having justified beliefs, collectively speaking, tends to lead to true beliefs (via the transfer of justification), this instrumental value that accrues to justified beliefs in virtue of their tendency to have desirable consequences itself depends on the point that (A4) justified beliefs themselves are likely to be true – and this latter relation between justified beliefs and true beliefs is not like the relation between means and the (desired) consequences they tend to produce; it is not a causal “productive” relation at all.

If one really wanted to say that the epistemic status of a belief as justified or unjustified derived from its power as a means to truth in the ordinary sense of “means,” that is, from the tendency of the belief to produce true beliefs, one would be in serious trouble. A belief might causally contribute to the acquisition of any number of false beliefs even though it is justified. Say you are about to watch a news-show because you believe it to be a good source of information: your belief may well be justified even though it will soon lead you to acquire lots of false beliefs through watching a show that is in fact a source of massive misinformation. Conversely, a belief might causally contribute to the acquisition of any number of true beliefs even though it is unjustified. Say you are about to watch a news-show because you believe it to be a good source of information: your belief may well be unjustified even if it will soon lead you to acquire lots of true beliefs through watching a show that is in fact a good source of information.

I should mention a problem that arises with respect to A4. On the one hand, A4 is highly plausible; on the other hand, it is unclear whether it can remain plausible while keeping the promise to connect justification to truth. The problem is this. What does it mean to say that a justified belief is “likely” to be true? Justified beliefs are based on grounds, typically referred to as evidence: mainly other beliefs, memories, perceptual experiences, and introspective experiences. To say that a justified belief is likely to be true can mean that its truth is likely *relative* to its grounds. This seems to fit with our concept of justification. Unfortunately, a belief that is likely relative to some grounds may be unlikely relative to other grounds, and may not be likely at all in any absolute sense. So, on this construal, it is not quite clear whether A4 really manages to connect the JB goal to the goal of having true beliefs. Alternatively, one might propose that justified beliefs are likely in some *absolute* sense, which would seem to secure the connection to the TB goal. Unfortunately, this may not fit well with our ordinary concept of justification. Suppose, as Descartes did, that a powerful evil demon leaves our perceptual experiences and memories just as they are but makes all the perceptual beliefs we base on them false. In this scenario, our perceptual beliefs would not be likely to be true in the absolute sense. So they are all counted as unjustified, even though they would still be based on the same grounds as our beliefs are actually based. This goes against widely shared, though not universal, intuitions about the concept of justification.

The problem is a tangled one, especially since it is not even clear whether it really is a problem. How much weight should really be given to intuitions about far-fetched scenarios like demon worlds? Why isn’t relative likelihood of justified beliefs enough to secure the connection between the JB goal and the TB goal? It turns out that these

issues are rather complex, and I will have to set them aside for now. It should be remembered, in any case, that they pertain to A4 rather than A1 and A3. To acknowledge that the plausibility of A1 rests largely on A3 is one thing. To uncover difficulties with the additional idea that the plausibility of A3 rests on A4 is another thing.

V The True-belief Goal and the Knowledge Goal

Given the considerations of the previous section, we can say that we (ought to) want to have justified beliefs because we want to have true beliefs, rather than the other way round. Let us conclude from this, at least as a working hypothesis, that the goal of having true beliefs is more basic than the goal of having justified beliefs; that the latter can be derived from the former. Even if this is correct, it does not follow that the goal of having true beliefs is the most basic epistemic goal. We still have to ask how it relates to the goal of having knowledge. Consider, then, the following claims:

- C1 If you want to have TBs you ought to have K.
- C2 We want to have K because we want to have TBs.
- D1 If you want to have K you ought to have TBs.
- D2 We want to have TBs because we want to have K.

If the C's were clearly more plausible than the D's, we could say that we (ought to) want to have justified beliefs *and* (ought to) want to have knowledge because we want to have true beliefs, but not the other way round. This would give us a reason for thinking that the TB goal is the basic epistemic goal.

However, there is something odd about the Cs. Consider C1. It is quite OK when taken collectively, as a claim about the general epistemic utility of having knowledge vis-à-vis the goal of having true beliefs. For knowledge itself constitutes evidence (Williamson, 2000): to the extent that you acquire beliefs properly based on knowledge you already possess, your new beliefs will tend to be true. But taken as a general claim about arbitrary individual propositions *p*, C1 seems deviant. The reason is not hard to find. Since knowledge requires true belief, part of what C1 tells you is that, if you want to believe *p* iff it is true, then you ought to believe *p* iff it is true.

It looks like the Ds are much more plausible. Consider D1. Quite independent of any instrumental value true beliefs might tend to have for acquiring knowledge (e.g., special true beliefs about how to acquire knowledge or evidence), D1 is plausible because true belief is a necessary condition for knowledge. It derives from a general claim about relativized goals with respect to individual propositions *p*: for any *p*, if you want to <know *p* if it is true, and believe *p* only if you know it> then you ought to <believe *p* if it is true, and believe *p* only if it is true>. On the face of it, it looks like we should say that we (ought to) want to have true beliefs because we want knowledge, rather than the other way round. This would give us a reason for thinking that the K goal is a more basic goal than the TB goal. We should, then, consider the competing hypothesis that the K goal is the basic epistemic goal: we (ought to) want justified beliefs because we want true beliefs, and (ought to) want true beliefs because we want knowledge.

The knowledge goal and the true-belief goal are both truth goals. So our two hypotheses are but two versions of the idea that truth is the primary epistemic goal.

Nevertheless, they would seem to be competing hypotheses. Most philosophers who have advanced or considered the idea that truth is the primary epistemic goal have taken this to mean that the true-belief goal is the primary one. Only a few – Williamson (2000) being the most recent – have explicitly advanced the view that it is rather the knowledge goal that is the primary one. Which view is the more plausible one?

Consider this thought: having true belief is a simpler, and in this sense, more basic state than having knowledge; having true belief is also a goal; hence, having true belief is a more basic goal than having knowledge. If “X is a more basic goal than Y” just means that X is somehow a more basic thing and is also a goal, then the true-belief goal may indeed be “more basic” than the knowledge goal. But this is not what one would normally mean when saying that X is a more basic goal than Y. Rather, what one would normally mean is that X is more basic *as a goal*; that we want, or ought to want, Y for the sake of X, rather than the other way round; or that the goal of having Y is in some other way derivable from the goal of having X. So, even if the state of having true belief is somehow more basic than the state of having knowledge, this does not show that it is more basic as a goal.

Let us try a different line of thought. The concept of knowledge is but the concept of justified true belief. So to want knowledge is not to want anything more than justified true beliefs. We have already seen that we want justified beliefs because we want true beliefs, rather than the other way round. Hence, we want justified true beliefs because we want true beliefs. The TB goal is the most basic goal after all. The impression that the K goal might be more basic than the TB goal is an illusion created by the complexity of the concept of knowledge.

This line of reasoning immediately faces the objection that, as Gettier (1963) shows, knowledge isn't merely justified true belief. Assume you believe that one of your friends owns a Corvette because you have adequate but nevertheless misleading evidence for the belief that your friend Smith owns a Corvette. Your friend Smith does not own a Corvette; but your friend Jones does. Then your belief that one of your friends owns a Corvette is both justified and true, but it isn't knowledge. Intuitively, your belief falls short of knowledge because it is only accidentally true.

Maybe we can make constructive use of this objection by turning it into a suggestion for how best to conceive of the goal of having true beliefs. Consider an initial diagnosis of Gettier cases: they show that knowledge requires *non-accidentally* true belief. Maybe this indicates that the true-belief goal is best conceived as the goal of having non-accidentally true belief. The goal might then be characterized by using subjunctives: we want to be such that we would believe p, if p were true, and would not believe p, if p were not true. Or, as Sosa (2003) might rather put it: it is the goal of being such that one would believe p if and only if p were true. (As a proposal for an analysis of the nature of knowledge that is responsive to Gettier cases the non-accidentality condition seems little more than a first step. However, it may well be good enough for paraphrasing our ordinary concept of knowledge and – more pertinent to the concerns of this paper – for describing the content of the true-belief goal). With the goal reconceived in this manner, the line of thought given above can be revived – now starting from the claim that to want knowledge is to want nothing more than justified non-accidentally true belief. As a result we get a defense of the idea that the true-belief goal is the primary epistemic goal in the slightly revised (improved?) version that the goal of having non-accidentally true beliefs is the primary epistemic goal.

VI Problems

The view that the goal of having true beliefs is the primary epistemic goal appears to be committed to the claim that we *ought* to want knowledge and justification *only* because, and insofar as, we want true beliefs. In other words, the view seems committed to a version of epistemic “value monism” (DePaul, 2000): the sole basic epistemic value, or good, is true belief; and the sole basic epistemic disvalue, or bad, is false belief. The other epistemic goods, knowledge and justified belief, are derived goods; they are valuable only because true belief is valuable; they are extrinsic epistemic goods, whereas true belief is the sole intrinsic good, epistemically speaking.

This view has to face Meno’s question: “It makes me wonder, Socrates, this being the case, why knowledge is prized far more highly than right opinion” (Plato, *Meno*, 97d). Meno has observed that intuition speaks in favor of value judgment (a) (below). One can make additional observations concerning intuitive judgments about epistemic value that will give rise to similar questions:

- (a) Knowledge seems better than mere true belief.
- (b) Justified true belief seems better than unjustified true belief.
- (c) Unjustified false belief seems worse than justified false belief.
- (d) Unjustified true belief versus justified false belief? Intuition hesitates.

If knowledge and justification derive their whole value from the epistemically basic value of true belief, then their presence shouldn’t add (and their absence shouldn’t subtract) any value from states already containing true belief or false belief. But then, why (a)? And why doesn’t intuition tell us that the states mentioned in (b) and (c) are equally good and bad respectively? And why does intuition waver about (d)? Why doesn’t it speak firmly in favor of true belief, however unjustified?

The value monist might ward off Meno’s original question, the one concerning (a), by citing the revised true-belief goal, the goal of having non-accidentally true belief: knowledge is better than mere true belief, because knowledge contains non-accidentally true belief. Of course, this will lead to another difficult question: Why is non-accidentally true belief better than mere true belief? (Timothy Williamson, who is not a true-belief monist, holds a view reminiscent of Plato’s own view: knowledge is better than mere true belief because it is more stable, less vulnerable to rational defeat by future evidence; cf. Plato’s *Meno*, 97–8, and Williamson, 2000, pp. 62, 78, 86). In any case, the revised (improved?) true-belief goal doesn’t seem to be of much help to the true-belief monist when it comes to the questions concerning (b), (c), and (d): these questions remain, even if the true-belief goal is reconceived.

Consider the other version of epistemic value monism, the one according to which knowledge is the sole basic epistemic good with true belief and justified belief as derived goods, valuable only because they are required for having knowledge. In a way, this view seems to fit with (a) to (d). Knowledge is better than mere true belief, because the whole value of the latter derives from the value of knowledge. Items (b) to (d) would be accounted for in terms of how many goods required for knowledge the states mentioned there contain or lack. But in another way, this version of monism seems deeply strange. How does knowledge get its basic value if true belief and justified belief don’t have any

basic value themselves? This seems like adding 0 to 0 to get 1. Does it get its basic value from the non-accidentality of true belief? But then that wouldn't be value *monism*: non-accidentally true belief would have to be a second basic epistemic value in addition to knowledge (cf. DePaul, 2002, p. 181). Or is the view simply that knowledge boils down to nothing more than non-accidentally true belief?

Now consider a pluralistic view about basic epistemic value. There are two basic epistemic goods, together with their accompanying bads: true belief and justified belief; false belief and unjustified belief. Knowledge, on this view, is a derived good, valuable only because it contains the basic goods. Evidently, this picture makes for the most natural fit with (a) to (d). Knowledge and justified true belief are better than mere true belief because they contain all the goods. Unjustified false belief is worst because it lacks all the goods. Intuition wavers about unjustified true belief versus justified false belief, because each contains one of the two goods and one of the two bads, and we can't decide, at least not right away, whether one of the two should weigh more heavily than the other.

Of course, on this pluralistic picture truth is not the primary epistemic good or goal after all. Instead, having justified beliefs is an additional epistemic value independent from true belief; it is at least as basic an epistemic good as having true beliefs. What, then, about the earlier considerations from section IV, supporting the view that the goal of having true beliefs is more basic than the goal of having justified beliefs? The pluralistic picture needs an alternative diagnosis of the plausibility of A1 – the claim that you ought to have justified beliefs if you want to have true beliefs. The pluralist, it seems, has to say this: A1 is indeed highly plausible; but not because of A3 and A4. Rather, A1 is plausible because it is trivial. You ought to have justified beliefs anyway, never mind whether you want to have true beliefs or want to have false beliefs or neither: you just ought to have justified beliefs, period. You ought to want justified beliefs for their own sake.

This alternative diagnosis of A1 is off the mark. It confuses epistemology with ethics (Firth, 1998a). If there is a proper place for absolute *oughts* – for *oughts* that have a hold on us no matter what – then their place must surely be in ethics. (Pointing out that the *ought* in question is the “epistemic” *ought* does not help; for this just leads to a reformulation of the very issue under discussion: What is the epistemic *ought*, if not the one that has a hold on us insofar as we want to possess truth?) Moreover, the position that justified belief is valuable for its own sake leads straight into a mystery: Why is it that, when we do attempt to justify our beliefs, we produce evidence we take to be indicative of their truth? Why do we regard something as evidence for a proposition, only if we take it to indicate that the proposition is true, or at least likely to be true?

So let us return once more to the idea that having true beliefs is the sole basic epistemic good and the primary epistemic goal. Bernard Williams (1978, pp. 37–46) attempts to give an account of how a “pure inquirer,” an inquirer who starts out with the basic desire for having true belief about some issue, would on reflection be led to want knowledge. Much adapted to our present framework, the account goes roughly like this. Assume you want true belief with respect to a certain question. Assume you are aware that (A1) if you want to have true beliefs, you ought to have justified beliefs. This gives you a reason for wanting justified belief on the issue in question. You now have two desires (“wants”): you want to have true belief on the issue, and you want to

have justified belief on the issue. If you are minimally reflective, you will notice this: having both your desires satisfied requires being in the state of having justified true belief on the issue in question. This gives you a derivative reason for having justified true belief on the issue, which gets us close to an explanation of why it is that you might want knowledge because you want true belief (C2). But wait, knowledge isn't merely justified true belief. Fair enough, so we should give a slightly revised account starting from the goal of having non-accidentally true belief.

The account offers the true-belief monist a fairly natural interpretation of C2: the claim that we want knowledge because we want true beliefs. If you want (non-accidentally) true belief and are aware of A1, then a desire for knowledge will tend to come along with the ride. The desire for knowledge turns out to be an offshoot of the desire for true belief.

Does the account also help the true-belief monist with the *Meno*-problems raised by the intuitions catalogued under (a) to (d)? Here is a proposal (see Sosa, 2002, for a different proposal). If you want to have true belief with respect to some question and are aware of A1, then you will have *two* desires: a desire for true belief on the issue and a desire for justified belief on the issue. Note that both are real desires. Even though one desire, the desire for justified belief, is directed at a derived good without basic epistemic value, it is nevertheless a real desire, just as real as the desire for true belief – desires for derived goods are no less real *as desires* than desires for basic goods. But we usually want our desires to be satisfied; and we prefer it if more of our desires are satisfied rather than fewer. This – I am inclined to propose – is the source of the intuitions catalogued under (a) to (d). The intuitions arise due to a confusion of sorts. They do not reflect any bonus of intrinsic value accruing to knowledge over and above (non-accidentally) true belief, nor do they reflect any intrinsic value accruing to justified belief that would be independent from the value of (non-accidentally) true belief; rather, they reflect our desire to have our desires satisfied.

References

- BonJour L. (1985) *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- DePaul, M. (2001) Value monism in epistemology. M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, ed. M. Steup. Oxford: Oxford University Press.
- Firth, R. (1998a) Are epistemic concepts reducible to ethical concepts? In *In Defense of Radical Empiricism: Essays and Lectures by Roderick Firth*, ed. J. Troyer. Lanham: Rowman & Littlefield.
- Firth, R. (1998b) Epistemic utility. In *In Defense of Radical Empiricism: Essays and Lectures by Roderick Firth*, ed. J. Troyer. Lanham, MD: Rowman & Littlefield.
- Gettier, E. (1963) Is justified true belief knowledge? *Analysis* 23: 121–123.
- Goldman, A. (1999) *Knowledge in a Social World*. Oxford: Clarendon Press.
- James, W. (1911) The will to believe. In *The Will to Believe and Other Essays in Popular Philosophy*. New York: David McKay.
- Plato (1976) *Meno*. Indianapolis: Hackett.
- Sosa, E. (2001) For the love of truth. In *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, ed. A. Fairweather and L. Zagzebski. Oxford: Oxford University Press.

- Sosa, E. (2003) The place of truth in epistemology. In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, ed. M. DePaul and L. Zagzebski. Oxford: Clarendon Press.
- Williams, B. (1978) *Descartes: The Project of Pure Inquiry*. London: Penguin.
- Williamson, T. (2000) *Knowledge and Its Limits*. Oxford: Oxford University Press.

Further reading

- Alston, W.P. (1989) Concepts of epistemic justification. In *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, NY: Cornell University Press.
- Alston, W.P. (1993) Epistemic desiderata. *Philosophy and Phenomenological Research* 53: 527–551.
- Chisholm, R.M. (1957) *Perceiving: A Philosophical Study*. Ithaca, NY: Cornell University Press.
- David, M. (2001) Truth as the epistemic goal. In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, ed. M. Steup. Oxford: Oxford University Press.
- Feldman, R. (1988) Epistemic obligations. *Philosophical Perspectives* 2: 235–256.
- Firth, R. (1998) Epistemic merit, intrinsic and instrumental. In *In Defense of Radical Empiricism: Essays and Lectures by Roderick Firth*, ed. J. Troyer. Lanham, MD: Rowman & Littlefield.
- Foley, R. (1993) *Working without a Net: A Study of Egocentric Epistemology*. Oxford: Oxford University Press.
- Fumerton, R. (2001) Epistemic justification and normativity. In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, ed. M. Steup. Oxford: Oxford University Press.
- Goldman, A. (1986) *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goldman, A. (1992) What is justified belief? In *Liaison: Philosophy Meets the Cognitive and Social Sciences*. Cambridge, MA: MIT Press.
- Haack, S. (2001) “The ethics of belief” reconsidered. In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, ed. M. Steup. Oxford: Oxford University Press.
- Kvanvig, J.L. (1998) Why should inquiring minds want to know? *Meno* problems and epistemological axiology. *Monist* 81: 426–451.
- Kelly, T. (2003) Epistemic rationality as instrumental rationality: a critique. *Philosophy and Phenomenological Research* 96: 612–640.
- Maitzen, S. (1995) Our errant epistemic aim. *Philosophy and Phenomenological Research* 55: 869–876.
- Nietzsche, F. (1989) *Beyond Good and Evil: Prelude to a Philosophy of the Future* (trans. W. Kaufmann). New York: Vintage Books (original work published 1886).
- Steup, M. (1996) *An Introduction to Contemporary Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- Velleman, D. (2000) On the aim of belief. In *The Possibility of Practical Reason*. Oxford: Oxford University Press.
- Wedgwood, R. (2002) The aim of belief. *Philosophical Perspectives* 16: 267–297.
- Williams, B. (1973) Deciding to believe. In *Problems of the Self*. Cambridge: Cambridge University Press.
- Williams, B. (2002) *Truth and Truthfulness: An Essay in Genealogy*. Princeton, NJ: Princeton University Press.
- Zagzebski, L.T. (2003) Intellectual motivation and the good of truth. In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, ed. M. DePaul and L. Zagzebski. Oxford: Clarendon Press.
- Zimmerman, M.J. (2002) Intrinsic vs extrinsic value. In *The Stanford Encyclopedia of Philosophy*, ed. E.N. Zalta. <http://plato.stanford.edu>.