

JUSTIFIED, TRUE BELIEF: IS IT RELEVANT TO KNOWLEDGE?

Jamie Leaser
Rutgers University

One who keeps herself up-to-date on recent epistemological trends may find at least two distinct camps currently at odds on the nature of knowledge. On the one hand, we find a group of epistemological iconoclasts who seem to favor significant conceptual change with regard to the concept 'knowledge,' and a group of traditionalists. I take the epistemological iconoclasts to be largely those reacting against the 'justified true belief' (hereafter JTB) account of knowledge on the grounds that if it captures any of what should be called knowledge, it only captures a small subset, namely 'knowledge that' or propositional knowledge, while neglecting skill knowledge and other nonpropositional cognitive achievements. I take the traditionalist camp to be comprised largely of those defending JTB (and hence, propositional) accounts of knowledge. In addition, some traditionalists try to accommodate their JTB accounts against threats that we may be 'Gettierized;' thus, such individuals take knowledge to be justified, true belief plus another condition to ward off allowing Gettier cases to count as instances of knowledge. These two camps appear largely antithetical to each other, at least at first glance. But both seem to hold that knowledge is a universal norm that is recommended for human pursuit and that provides a standard against which actual human cognitive achievements are evaluated. Where they differ is most apparent in the areas of how determining the universal norm or standard should be conceived and how well human cognitive achievements 'measure up' against this standard.

Traditional epistemologists tend first to establish the criteria for knowledge, and following that, to evaluate human achievements against their standard. Some recent epistemologists guide their choice of criteria by first selecting some of our highest cognitive achievements, or apparent paradigm instances of knowledge, and then deriving a conception or theory based on paradigm instances of how we use the term 'knowledge.' Thus, presently I want to alternate between viewing how we ordinarily use the word 'knowledge' and how I think we ought to use the word 'knowledge.' From this alternation I wish to carve a rough conception of knowledge by exploring several questions. But note, to begin, that my overall motivation is to operate according to a *principle of preservation*; I want to preserve as much of the traditional (JTB) account of knowledge as possible, while regarding cognitive endeavors that may not properly be called 'knowledge' as equally intellectually and

pragmatically valuable. Some would argue that our use of the term 'knowledge' should expand to become applicable to various nonpropositional cognitive achievements, e.g., skills. Thus, I first want to attempt to clarify what we mean by knowledge embodied in skills. In clarifying, I will ask whether we ought to apply the term 'knowledge' to the cognitive efficacy embodied in skills and if so, whether such an application serves to undermine the conception of knowledge as justified true belief.

Second, I will discuss the following question. If we conceive of knowledge roughly as justified true belief, and since scientific beliefs often turn out to be false, then is science not the knowledge generator we think it is? But since we take science to be the area in which our greatest cognitive achievements occur, claiming that science does not produce the vast body of knowledge we think it does calls into question a paradigm instance of knowledge. Thus, I will address whether we should adjust our concept of knowledge to accommodate science as a paradigm instance, or preserve the traditional concept of knowledge and reject the notion that science presents a paradigm instance of knowledge.

In answering my first and second questions, I will defend the JTB account of knowledge, although my doing so will be rife with qualifications and compromises. My third aim is to sketch an account of justification I find most appropriate if one is to preserve the JTB definition of knowledge, namely, one that allows the criteria for being justified to vary according to the individual's context and goals, thereby possibly making justification and knowledge more accessible than other JTB accounts. Finally, I will conclude that we ought to preserve the concept of knowledge as roughly that of justified true belief, but that our set of cognitive values and ideals must expand to include more than just knowledge. Rather than change our concept of knowledge, perhaps we ought simply to allow other intellectual endeavors to share the cognitive limelight; that way, we can avoid undergoing a bewildering conceptual change and yet also avoid shortchanging more recent cognitive achievements that receive growing attention. In addressing the above concerns, I will begin by discussing the most recent proposed inclusion to knowledge, namely, 'knowledge how' or knowledge embodied in skills.

I. Skill Knowledge: What it is and What it Means for JTB Theories

Epistemologists concern themselves with skill knowledge both as a result of the work of Michael Polanyi (1958, 1967) and perhaps as a result of political issues aiming to address whether 'knowledge

how' has been excluded from what has traditionally been called knowledge because of its association with the working class, manual laborers and craftsmen. Some might actually find skill knowledge more basic than propositional knowledge. While there may be good reasons for taking skill knowledge to be more basic, one first needs to clarify the notion of a skill and how it embodies knowledge. Since I do not know of any necessary and sufficient conditions for defining a skill, and since we refer to a skill as 'knowledge how', we might first look at how we use the phrase 'knowing how'. We say of young children that they know how to feed themselves, tie their shoes, tell time, and so on, while we say of teenagers and adults that they know how to play the piano, drive, ride a bicycle, swim, sew, fix a car, knit, and so on. These activities may have little in common other than that they are all learned at some point, and that they are learned in stages. Learning in stages seems central to a skill to differentiate it from activities that one tries and either can or cannot do from the first try, for example, curling one's tongue or contorting one's limbs in a 'doublejointed' position; we typically say that one *can* curl her tongue but not that she *knows how* to curl her tongue since such an ability is genetic. In addition, Brown (1992) claims that skills are learned by training and practice, improve with practice, deteriorate when not exercised, and are exercised fallibly even among those with great expertise. However, and as is to be expected, one also finds scores of borderline cases, activities that one learns but learns quickly, such as skipping or snapping one's fingers. Are such activities skills? I am inclined to deny that they are because it seems that once one 'catches on', there is no more to learn. More likely, we will call those activities skills which not all adults can do, are typically learned in stages, show some kind of proficiency or expertise and reveal some sort of technique as they are being carried out.

The problem with defining a skill *en route* to deciding whether there is such a thing as nonpropositional skill knowledge is that we can, *ad hoc* of course, call all and only those things skills that either are or are not conceivable in propositional stages. Then, if I call only those activities 'skills' which I can reduce to propositional stages, I can assert that skill knowledge really is 'knowledge that' in clever disguise and still maintain that whatever is nonpropositional is neither a skill nor properly called knowledge. Our intuitions seem to be of little help since they are likely to be guided by our theories, so we may expect disagreement on which activities to call skills and embodiments of knowledge. Also, we want to decide whether such activities count as instances of knowledge to judge whether they are nonpropositional and whether the JTB account of knowledge will be rendered irrelevant if it cannot apply to such activities.

We can, of course, make headway even without a clear conception of what a skill is, as long as we use examples with the awareness that we are presupposing a notion of skill that is open to objections. Also, with regard to the question of whether the knowledge embodied in skills is propositional, in examining performed skills in which the individual cannot give a propositional account of what she is doing, we must consider reasons why she cannot give a propositional account of what she is doing other than the reason that *no such account can be given*. Recall that Harman (1994) presents an unusual view of foundationalism in which one is only justified so long as she can keep track of the justificatory chain supporting her belief. According to such a version of foundationalism, if I cannot reproduce the steps in Goedel's proof I am no longer justified in believing the conclusion, even though I could reproduce it if I had adequate time to check my logic book and 'refresh my memory' as to how I learned it in the first place. A more natural response might be to assume that I am still justified in believing the Goedel proof since I remember in a more general sense how I came to believe it although I have since forgotten the specific steps by which the conclusion is obtained. Propositional steps were an indispensable part of the learning process, although I cannot now tell someone in a satisfactory way how I learned the proof.

A similar process to the learning of a proof might take place with many skills. A series of propositions are learned to acquire the skill, but once the skill can be successfully performed, the propositions gradually fall out of memory. But note, I am not suggesting that learning a series of propositions is *sufficient* for learning a skill, nor that learning a skill is *algorithmic*. I am merely postulating that learning a series of propositions may be necessary for learning a skill, and that the successful performance of the skill may be similar to the conclusion of the proof; once the end result is reached, the steps are forgotten until one is called upon to reproduce them. Again, I am trying to offer *suggestion's* as to why 'knowledge how' cannot be explained propositionally (by the one performing the skill) in many cases without asserting that this must be (or must not be) the case. Rather, I want to propose possible explanations for why the JTB account of knowledge (and hence propositional knowledge) may still be involved in 'knowing how.'

Consider swimming as an example of a skill.¹ Talented competitive swimmers usually learn how to swim at an early age. (Of

¹ Again, my selection of a particular activity as a skill is tentative. If one disagrees that swimming is a skill, one will probably also disagree with the conclusions I draw about skill knowledge from the example of swimming. Thus, my minimal hope is that if one agrees that swimming is a skill, she will agree with my conclusions drawn from the example.

course, we may repeat an analogous statement for talented athletes and musicians in general.) Such talented swimmers probably do not remember much about what it was like when they learned how to swim, and they may not be very good at teaching others how to swim. On the other hand, they may be very good at teaching other talented swimmers how to improve their technique. Perhaps the more talented swimmer can frequently explain to the fellow talented swimmer how to improve her technique because it is less difficult to think of explanations in terms each would understand, which may not be the case with the beginning swimmer. Explaining the acquisition of a skill to one who attempts to learn it for the first time may require the expert to refrain from employing skill jargon and instead employ simpler terms in the explanation, which may tax her memory. She may experience a tax on her memory simply because she is not frequently called upon to offer such 'layperson' explanations.

Similarly, graduate students in philosophy who converse solely with fellow graduate students may have difficulty explaining philosophical notions to introductory students without employing philosophical jargon which the student does not yet understand. Graduate students may have even greater difficulty doing so than professors, because graduate students simply may not yet be accustomed to offering such explanations and (employing different terms) to those of varying degrees of philosophical background and competence, while professors, though much more philosophically advanced than graduate students, have had more practice at giving such explanations. In short, one reason those with skill knowledge may not be able to explain how they perform their skills may be a simple communication barrier; they may be able to explain much of what they do to someone slightly less knowledgeable (but still versed in the jargon of the skill), but not to someone who does not even understand many of the concepts required for a simple explanation.

The second reason one may not be able to explain the acquisition or performance of her skill is that she cannot remember what it was like to learn the skill at that level. We see examples of this problem in attempts to explain how to write well to a student whose paper has an abominable number of simple grammatical errors. I may remember how I learned to improve my writing style in various stages in college, but not in grammar school. Such a stage in my learning process has long been forgotten. Thus, I do not remember how I learned to avoid making the type of errors made by the student I am required to help, but if I visited my former grammar teachers I probably could get a general idea of how I corrected my grammatical misconceptions and progressed at that level and then be of greater help to the struggling student. I may have had to learn a series of propositions to progress to the point I am at today, but the

propositions that I was required to know to learn that skill are no longer part of my working memory. In some cases, (with swimming for example), at each stage of skill acquisition or improvement I learn new propositions which may assume knowledge of an earlier stage. But what I remember about the earlier stage will be only enough to preserve my existing knowledge; I will not remember all that was taught to me in order to explain how to do something "from scratch" at an earlier stage.

The main point I want to make from the above discussion is that we have to be wary of hastily concluding that (at least partially) propositional accounts cannot be given of skills just because when we ask those engaged in the skills they cannot immediately offer a propositional account of what they are doing. It seems that propositions are involved in the skill even if propositions cannot provide an exhaustive account of what is going on. Evidence for the notion that propositions are at work in performing skills may be found when we watch someone practice her skill and learn how to improve. Practice sessions often involve breaking down the skill into various steps and analyzing those steps so that one learns how to perform the skill more efficiently. Improving one's technique at a skill such as swimming, for example, is highly dependent on propositional knowledge. Factual knowledge, such as an understanding that the greater water resistance one causes, the more slowly she will move, enables one to strive for an aquadynamic technique. Swimmers are taught how to move their hands in an \sim shape while pulling to move the greatest amount of "still water" and achieve the maximum efficiency out of each stroke. In short, if the goal is to move as quickly as possible through the water, a talented swimmer usually holds many justified, true beliefs about why moving her arms and legs in a particular way will expend the least amount of energy and have the most efficient stroke technique, thereby achieving her goal.

Also, some activities may involve repetitive motor skills that are mere physical habits that can occur apart from any particular cognitive state, just as one may drum her fingers or jiggle her foot without even realizing she is doing so. It is not as though such activities involve nonpropositional knowledge; rather, some appear to be almost noncognitive. But even in repetitive skills, one apparently has a belief that she ought to carry out the particular task in the same way she did previously, because the way in which she carried out the task previously enabled her to achieve her goal. In short, if we think of carrying out a skill as intentional action, then we will assume that it involves beliefs and desires of some sort. Presumably, one has at least a belief about what she is doing that helps to explain the repetitiveness, namely, that the way in which she

carried out the task before enabled her to achieve her goal and carrying out the task the same way in the future will continue to enable her to achieve her goal. For if one holds no beliefs about whether she is carrying out a task appropriately, the action seems unintentional, and certainly not properly embodying knowledge, but merely an accidental maneuvering, or with some skills, beginner's luck.

The JTB theorist may also allege that if one performing a skill holds no beliefs about what she is doing while she is doing it, it is unclear how knowledge is embodied in the activity. For knowledge presumably cannot apply to motor activity; it typically applies to cognitions. But one's cognitions either are or are not about one's motor activity during the performance of the skill. If they are not at all, then the JTB theorist may fail to see how knowledge is in any way embodied in what takes place while that skill is being carried out. If they are, then it seems that propositional knowledge is bound to come back into the central picture.

I am not trying to suggest that skill knowledge can be entirely grasped in propositional language or is even convertible into propositional language. Let us consider another example. Gymnasts often report having a 'feel' during a routine or a dismount in which they are still in the air but claim to know how they are going to land based on how they "took off." They may very well know how they will land, or be very accurate predictors of their landings while they are still in the air. How do we account for this 'feel' that seems a part of their skill knowledge? I doubt that gymnasts possess a body of propositions that they recite to themselves at this point; rather, it seems as if they take a kind of mental picture or snapshot of how they felt overall once before while successfully dismounting and remember it as an instance of what the achievement of the goal felt like. They may not remember exactly what they did in the process of the last successful dismount but may recall by trial and error that a particular feeling during the "take-off" is familiar as an instance of success or failure. Beliefs that one's mental picture of the event match a previous successful dismount can still be true or false; one can, of course, think she will have a successful dismount based on her belief that her overall feeling matches a feeling she had when she successfully dismounted, but still be wrong. Also, such mental pictures of what the right dismount felt like may be combined with some simple "physics" beliefs that also allow her to predict her landing. But her beliefs and desires about what she is doing while she is doing it are what differentiate her from the individual who simply dismounts successfully by accident, never having performed a dismount before and never having thought about how it should work.

I wish to draw several points from the preceding discussion of skill knowledge. First, we may conceive of a skill roughly as an activity which is not typically performed by all adults, is learned and improved in stages, and is performed with varying degrees of expertise and with varying techniques. Second, there is a valuable cognitive achievement involved in the acquisition, performance, and perfection of a skill. One exercising a skill has a recognized goal which can be achieved more or less successfully, and we may evaluate her efficacy and expertise accordingly. Presumably, such an evaluation also requires a norm or standard against which skill performance is measured. Third, one's success at exercising a skill is usually manifested by a complex combination of physical dexterity or agility and cognitions including, (but not exhausted by), beliefs and desires. Furthermore, I take the beliefs and desires involved in the successful exercising of a skill generally to be true and justified. What does this mean for the JTB definition of knowledge? It does not seem that JTB accounts of knowledge are undermined simply by showing that what we call skill knowledge is not wholly propositional. Because skill knowledge is not wholly *nonpropositional*, JTB accounts still seem relevant to explaining the nature of skill knowledge. The challenge I see facing JTB accounts of knowledge with regard to skill knowledge is not whether an exhaustive explanation can be given in propositional language, but whether sufficient sense can be made of skill knowledge in propositional language.

Also, because I know of no alternative conceptions of knowledge that account more adequately for skill knowledge and possess the overall explanatory power that JTB accounts have, at this point I favor preserving JTB accounts of knowledge and attempting to translate the murky processes at work in skill performance into propositional language where possible. Furthermore, although I hold that JTB accounts of knowledge should explain the exercising of skills where possible, it does not follow that I maintain that what we commonly refer to as 'skill knowledge' ought properly to be called knowledge, for it need not, and perhaps should not. Thus, two possible implications for JTB accounts are that (1) JTB theorists may claim that they provide an adequate account of skill knowledge insofar as *what is embodied in skills is properly called knowledge*, but also that (2) they need *not* say that *most of what is embodied in skills is properly called knowledge*. Rather, they may say that whatever cognitive efficacy or achievement is embodied in skills is equally *as valuable as* knowledge, but for the sake of conceptual simplicity, should not actually be called knowledge. Again, we may value the successful or expert exercising of a skill as highly as knowledge without actually including it under the rubric of

knowledge. But it is my contention that we ought to be more concerned with having the right overall cognitive values and ideals than with valuing only knowledge and then including as much as we possibly can within the concept of knowledge. In sum, the existence of nonpropositional skill expertise does not warrant our rejection of the JTB definition of knowledge.

II. The Problem of Scientific Knowledge

Another statement used to level a challenge against JTB accounts of knowledge is that we take ourselves to be in possession of a vast body of scientific knowledge. And yet, if knowledge is defined as justified, true belief, and given that most of our scientific beliefs turn out to be false, we do not in fact have the vast body of scientific knowledge that we think. So we either have to admit that our paradigm case of knowledge turns out not to be knowledge after all, or throw out the truth requirement for knowledge, thereby altering our concept to fit our paradigm case of knowledge. I take the former to be the more desirable alternative.

Relinquishing the truth requirement is a bit frightening primarily because we do not know whether another criterion will have to replace truth as necessary for knowledge, or whether something like mere justified belief will suffice. Because I know of no replacements for truth as a necessary criterion, I will only discuss the latter alternative. Suppose we define knowledge as mere justified belief. Presumably, we can still retain the notion that once one becomes aware that her belief is false, she is no longer justified in holding it. At least given this, truth is not entirely out of the picture because awareness of falsity undermines one's justification. If we do not require truth for knowledge, one who has a justified belief can already claim to possess knowledge, and since knowledge is our highest epistemic ideal, then should the individual find her belief is true,² nothing much is added to her epistemic situation. Perhaps the individual would have knowledge in a stronger sense, but no epistemic threshold would be crossed. Also, relinquishing truth as a necessary criterion for knowledge seems as though it could harm the practice of science, for if two scientific practitioners held justified beliefs which contradicted each other, one could assert that they both know, even though one claims to know A and the other claims to know \sim A. But then why would further scientific investigation be needed? Typically, if two individuals claim to know propositions which contradict one another, we assume that they cannot both know, and yet, if knowledge does not require truth, they can both

² I am speaking hypothetically here; I realize that one might charge that we never conclusively find out a belief is true.

know, which leaves us to wonder why further hypothesis-testing would be needed.

But what alternative do we have? Must we claim that scientists do not possess the knowledge we think they do? The answer is both yes and no. If we answer yes, again, such an answer need not demean scientific achievements in any way, for although knowledge is highly valued, we do not *merely* value the truth of the 'end state' concluding belief. We also value the methods and reasoning processes that give rise to the conclusion and the justification for that belief. If the processes giving rise to the conclusion are useful and frequently produce beliefs that seem likely to be true, we have little reason to despair. We respect scientific practice because it employs methods, reasoning, and testing procedures that seem to make arriving at truth (and knowledge) likely. But, one might object, if scientists in fact turn out to be wrong so much of the time, why should we think that these methods are in fact valuable and likely to produce knowledge?

We should continue to value the methods because *they are what tell us* we were wrong; they lead us to falsifications of previously held beliefs, and when we appear to have found out an hypothesis is in fact wrong, we still increase our body of knowledge, just not in the way that we expected to do so. This constitutes the 'no' part of my answer to whether we must claim that scientists do not possess the knowledge we think they do. Although theories often turn out to be false, scientists still compile a great body of particular true beliefs in the process of theory construction and testing. It seems that each time we find out we were wrong, we appear to gain an unexpectedly justified true belief. Admittedly, the new belief may not be as interesting or useful in terms of its predictive and explanatory power, but we still have the rigorous hypothesis testing and attempts at falsification that occur in scientific practice and allow the body of scientific knowledge to expand at a rate envied by practitioners in other fields. Such procedures are also what prompt us to praise scientific practice as the apotheosis of a rational endeavor. In sum, every falsification engenders a new truth or a new hypothesis, and what we find equally valuable in science are the methods of testing and attempts at falsification, so we need not bemoan any sudden evaporation of the vast body of 'scientific knowledge' even if we accept a JTB account of knowledge.

Thus far I have lobbied for the JTB account of knowledge largely from a defensive position. Next, I intend to sketch (from a more offensive position) how the JTB account might be cashed out. In particular, (as mentioned earlier) the two camps in epistemology differ as to how human achievements measure up against the standard of knowledge, and traditionalists are sometimes accused of making justification (and hence knowledge) unattainable. For this

reason I wish to propose an account of justification that makes knowledge more easily attainable, and is also not immune to influence by nonepistemic concerns.

III. A New Look at Justification

Justification can be earned in degrees, and so, as a component of knowledge, it seems that knowledge also can be earned in degrees. At the very least, we might claim to know in a weak or strong sense (Malcolm 1993). But we can also say that the standard of justification varies depending on the context or individual goals and perspective of the cognizer. Whether or not one is justified in believing *X* ought to depend in part on what one's goal is *with respect to believing X*, and what consequences (both epistemic and nonepistemic) might ensue as a result of believing *X*.

For example, Clifford (1993, p. 505) has argued that "it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence." When Clifford condemns those who believe upon insufficient evidence, his condemnation is both epistemic and moral, insofar as neglecting one's epistemic duties can be morally blameworthy when the lives of others rest on the truth of one's belief, and relaxing one's epistemic standards can make one epistemically lazy such that when serious consequences do rest on the truth of one's belief, one may not be able to subject the belief to appropriate, careful scrutiny because he or she has formed bad epistemic habits. Clifford offers the famous example of the ship owner who sends an old ship to sea without having it inspected, ignores possible evidence that the ship is unfit for sea travel, and instead merely trusts God to keep those aboard the ship safe. Clifford claims that the ship owner is both epistemically and morally guilty, although he implies that epistemic guilt is moral guilt in light of the consequences that might ensue from holding a belief without sufficient evidence.

Although Clifford raises an interesting idea by conflating moral and epistemic responsibilities, he seems to develop his idea in an overly stringent way, for most of us assume that hastily adopting an apparently inconsequential belief will not immediately allow one to slip into bad epistemic habits, nor should it constitute a grievous moral wrongdoing. It does seem important, however, for one to be aware of 'what is riding on' her belief. But if we allow someone to believe hastily when nothing much is at stake, will poor epistemic habits abound? I doubt that this has to be the case. We may allow for a kind of 'sliding scale' justificatory requirement given the individual goal of the cognizer without necessarily being threatened with epistemic pandemonium. There is a constant standard operative in

the assessing of one's goal, perspective, and reasonably foreseeable consequences in holding the belief. Thus, although the requirements for being justified with respect to a certain belief vary in different situations, the *method* of determining the requirements does not vary.

Annis (1993) offers the example of two individuals possessing the general information that polio is caused by a virus, claiming that the "issue-context" of the individual determines whether either has enough justification for knowledge. One individual is a non-medically trained person while the other is preparing for an examination for the M.D. degree. Annis maintains that the context of the individual in an M.D. examination requires that he have a great deal more information about polio (than the non-medically trained person) to qualify as being justified and possessing the knowledge that polio is caused by a virus. Annis sees the issue-context as "what specific issue involving *h* [the belief] is being raised. It determines the level of understanding and knowledge that *S* must exhibit, and it determines an appropriate objector-group." (1993, p. 282) The "objector-group" can basically be understood as the community to which *S* submits his belief for critical scrutiny. The objector-group is one that is informed about issues involving *h* and whose objections to *h* must be seriously considered by *S* in order for *S* to be justified in believing *h*.

We may also conceive of a context in terms of an individual's goal, perspective,³ and the foreseeable consequences of his belief. Thus, we might say that the non-medically trained person's goal in knowing whether polio is caused by a virus is simply to find out whether *he ought to avoid certain kinds of contact with those afflicted with polio*. The medically trained individual, however, desires to know for the reason that the first individual noted, but also to give others medical advice about the disease and to be considered an expert. Thus, the medically trained individual's knowledge will affect both the knowledge and the health of many other people. It appears, then, that his goal in knowing whether polio is caused by a virus requires that he have greater justification in order to possess knowledge. But what about the first individual? Will he form bad epistemic habits by allowing a softer standard for himself in this case? He need not, mostly because he is still required to discern his goal in knowing the proposition, his initial perspective, the possible consequences of holding the belief, and then judge how much justification is required of him in this particular case. Thus, one's

³ I am using 'perspective' to mean one's standpoint or social situation that reflects certain biases and also determines the appropriate 'level' (e.g., child, adult, expert, layperson) of objections. Awareness of one's initial perspective on the issue reveals one's biases that need to be brought to light and corrected or balanced as much as possible.

discernment of his goal, perspective, and foreseeable consequences and then assessing the amount of justification required in light of these three things will proceed in much the same way each time a belief comes under scrutiny; it is only the actual amount of justification required for knowledge that varies from instance to instance.

Also, presumably if one revises her goal or discovers that more serious consequences follow from her being mistaken than she previously thought, the individual needs to adjust her level of justification accordingly and, if possible, avoid bringing about such consequences before she has once again achieved the level of justification required. In one sense this may sound strange, but it actually fits with the way we behave in everyday life. For instance, we frequently mention to a friend in an offhand manner a belief or a piece of information for which we are minimally justified, only to have another individual suddenly join in the conversation and question what we just said. Perhaps our remark has particular relevance to the third individual's life, for example, a remark about a school to which the individual was considering sending her daughter. Upon hearing the remark, let us suppose the individual jumps in the conversation and wants to hear more, suddenly having doubts about sending her daughter to the school being discussed. At this point, we will usually qualify our remarks, reporting the source from which we heard the information, and perhaps adding "...but don't take my word for it!" It is this everyday phenomenon which I am attempting to capture in my proposal for a requirement for revising one's standard of justification upon discovering a new goal or consequence. Revision of a justificatory requirement seems appropriately understood in terms of one's goals and possible consequences of the belief, and is called for in the above example because suddenly "more is riding on" the belief when previously unforeseen consequences are introduced with the third interlocutor.

Admittedly, I see the amount of justification required for knowledge as relative to the cognizer. And, of course, because justification is context-relative, knowledge is as well, to some extent. My allowance of the relativity of justification and knowledge may cause some to question whether I am suggesting that truth might be variable in the same way. For example, am I secretly suggesting that the statement 'This surface is flat' is true given some goals, but not given others? I do not think so, although we surely have to allow that words mean different things in different contexts, but as long as we specify our meaning in such cases, the appearance of relativizing truth can be avoided. We do a sort of informal translation of terms quite frequently in ordinary language, and that is what is needed in the cases of statements that appear to be true given some goals, but

not given others. For example, nonphilosophers use the words 'proof' and 'valid' in a much more casual way than philosophers, and we often hear nonphilosophers make comments such as "That's a valid point." Are points valid given some goals, but not given others? I think it is simpler just to say that nonphilosophers use 'valid' to mean 'has some presumption in its favor' or 'pretty likely'. Also, some nonphilosophers use statements of the form 'Not all X are Y' and 'All X are not Y' interchangeably. If we fail to translate what such individuals 'really mean' in these cases, it looks as though we will be forced to admit that the statement 'All Germans were not Nazis' is true given the nonphilosopher's goals. Thus, I think we will have less confusion if we simply specify in such cases how a word or phrase is being used differently or replace the word or phrase with a more appropriate one.

Next, some might object to my conflating epistemic and nonepistemic goals and consequences relevant to belief-formation since I allow that epistemic requirements may be influenced by some pragmatic consequences of our holding a particular belief, or of our being wrong about the truth of a particular belief. However, doing so gives us a more flexible notion of justification and may facilitate our attainment of knowledge. Also, it is very natural at times to assess people epistemically while being influenced by pragmatic goals. For instance, let us suppose that we desire to know which plants are safe for humans to eat.⁴ Assume that individual edible plant-finders A and B have the same success rate; out of every ten plants they encounter, they determine eight out of the ten correctly (whether they are edible or poisonous). If we merely examine the number of true beliefs and false beliefs in each case, the number will be the same for each individual and they might be judged as equals with respect to the possession of knowledge about what edible plants look like. Indeed, if our goal is simply the tallying of true beliefs as opposed to false beliefs about edible plants, A and B appear equally successful at the task.

But suppose that A's errors always amount to believing edible plants to be poisonous, while B's errors always amount to believing poisonous plants to be edible. Will we still judge them as equals with respect to knowledge in spotting edible plants? I am inclined to think not, because when considering the goal of survival, and if one does not have the cognitive resources to be absolutely certain about which plants are edible, it is surely better to err in one direction than the other.⁵ Thus, in considering pragmatic goals, I am not advocating

⁴ This example is similar to one used by Stich (1990).

⁵ Note that one could not opt simply to avoid fatal *error* (ingesting a poisonous plant) by always judging the plants to be poisonous, for then one would starve and wind up committing another fatal error.

that pragmatic goals override epistemic goals, but rather that pragmatic goals can at least tell us in which direction it is better to err. We may say that the individual who errs in the direction that leads to our death by poisoning is more unjustified than the one who erred in the opposite direction. In short, pragmatic goals and consequences occasionally factor into epistemic evaluation and requirements, and an appropriate account of justification should be sensitive to this phenomenon.

I have mentioned that the individual has to determine her own goal and assess the level of justification needed to know in light of her goal, but how does this work? Can one be justified that she has ascertained the correct level of justification and the foreseeable consequences of the belief? Ultimately, this process is up to the individual; how each individual assesses the level of justification required for her to know in a particular case will slightly vary, and of course some individuals will be more judicious at such a task than others. I recommend Foley's (1993) suggestion of how we are to make up our own minds as applicable to how we decide how much justification to require of ourselves and how to decide when we have reached the appropriate level of justification:

I am to make up my own mind by marshalling my intellectual resources in a way that conforms to my own deepest epistemic standards. If I conduct my inquiries in such a way that I would not be critical of the resulting beliefs even if I were to be deeply reflective, then these beliefs are rational for me in an important sense, an egocentric sense. There are various ways of trying to spell out exactly what this amounts to, but for purposes here the details can be left open. The basic idea is that if I am to be egocentrically rational, I must not have internal reasons for retraction, ones whose force I myself would acknowledge were I to be sufficiently reflective. (1993, p. 148)

Foley goes on to clarify that he does not mean to imply by "deepest epistemic standards" that such a process of making up one's mind is some kind of private Cartesian process; rather, the process of deep reflection should involve holding open a belief to critical scrutiny from others (for example, Annis' objector-group), weighing counterevidence, and so on. And again, "marshalling intellectual resources" will be influenced by certain pragmatic considerations in deciding those cases in which we should continue to gather evidence, seek counterexamples, and address problems raised by the

objector-group. We may employ a Foley-style approach to one's *ascertaining* the degree of justification required that, when added to other necessary criteria, produces knowledge in a given situation. Of course, we can never be perfect predictors of the consequences that might ensue from our holding to a certain belief, but we can be deeply reflective and at least attempt to avoid epistemic self-blame even if others ultimately judge us to be epistemically blameworthy in a situation.

IV. Concluding Remarks

In closing, I have defended the thesis that knowledge ought to be defined as justified true belief in some sense. Primarily, I have argued that justification, truth, and belief should not be thrown out as necessary criteria for knowledge in light of recent charges that the JTB accounts can only offer explanations of a subset of knowledge, namely propositional knowledge, and say nothing informative about skills, which may comprise another subset of knowledge. I have not addressed the sufficiency of such criteria, particularly in view of the threat of being Gettierized, nor will I, primarily in the interest of space. But also, the insufficiency of justification, truth, and belief for knowledge does not warrant a *rejection* of the JTB account of knowledge; it merely warrants an *addition*. Because those who take the Gettier problem to be important usually look for a fourth criterion to add to justification, truth, and belief, I hope that saving a discussion of Gettier concerns for another time does not detract terribly from the present paper. I do not see the Gettier problem as a frontal attack on the JTB thesis, and so I do not see a discussion of the Gettier problem as seriously affecting what I have said about JTB accounts, except that another condition might have to be added.

At any rate, from my earlier discussion I want to conclude that justification, truth, and belief still are necessary for knowledge, even though apparently in some cases the individual may judge herself to require no justification for a particular belief.⁶ Allowing a kind of sliding scale justification relative to the individual's goals provides some flexibility that JTB accounts might traditionally have been accused of lacking, avoids full blown relativity, and accommodates many of our everyday intuitions about knowledge. Furthermore, the JTB account appears able to explain some of what is called skill

⁶ This can create somewhat of a word game, for one may say that the analysis and judgment that one needs no justification for a belief requires a reasoning process that is itself a justification for why that belief needs no justification, and I would be inclined to agree. Justification can easily though furtively be maintained as a necessary criterion for knowledge since one person's unjustified belief is another person's self-justified belief.

knowledge, albeit with some translation of what goes on in the mind of one performing the skill. However, I also maintain that we should only preserve the JTB account of knowledge if we recognize that knowledge does not stand alone as a cognitive value or ideal. In addition, knowledge need not be the only norm recommended for human pursuit, nor the only standard against which human cognitive achievements are to be evaluated. Obviously, if other nonpropositional cognitive achievements are intellectually valuable but are not properly called knowledge, then we will require other standards against which such achievements are to be evaluated. But rather than reject the JTB account of knowledge as insignificant if it only accounts for propositional knowledge and not other subsets of knowledge such as skills or 'knowledge how,' I want to preserve the traditional notion that knowledge is largely propositional insofar as it is an attempt to match a belief to a 'fact of the matter,' but add that knowledge is not the only cognitive state of any genuine value. Although I have not attempted to spell out all that is of cognitive value, we may at least hold that success or efficacy at our other, nonpropositional cognitive endeavors (e.g., skills) are to be valued as highly as knowledge. In this way, we avoid demeaning those cognitive activities that are not properly called 'knowledge' but are equally valuable and necessary to our survival.

In light of the above comments, some might question whether I really have avoided conceptual change. If we think of the role that 'knowledge' has played in our system of epistemic concepts, part of this role has been as the maximally valuable cognitive state and often as the *only* valuable cognitive state. Admittedly, I am accepting major changes in how we think about our epistemic condition, if indeed it is true that we previously thought of knowledge as the only valuable cognitive state. But merely allowing changes in how we think about our epistemic condition and our cognitive values does not, as I see it, engender the degree of confusion that arises as a result of conceptual change. Thus, I do not see accepting a shuffling of our cognitive values as a closet concession on my part.

Again, my main reason for defending justified true belief as knowledge and adjusting our epistemic value system has been motivated by a basic principle of preservation. I want to avoid conceptual upheaval where possible, and if we throw out justification, truth, and belief as criteria for knowledge to expand our concept of knowledge, we will (as I see it) either have to devise different criteria for various subsets of knowledge, or start 'from scratch' and devise new criteria for knowledge as a whole, still having to account both for propositional knowledge and for our efficacy in nonpropositional cognitive endeavors. On the other hand, if we hold on to our traditional criteria for knowledge, we set fewer

tasks for ourselves and instead expand our set of cognitive values. In doing so, we can (at least for the time being) avoid conceptual revolution and hopefully satisfy the interests of a few members from both the traditional and the iconoclastic epistemological camps.⁷

References

- Annis, D., (1993) 'A contextual theory of epistemic justification', in L. Pojman (ed), *The Theory of Knowledge: Classic and Contemporary Readings*, Belmont, Wadsworth Inc., pp. 280-287.
- Brown, H., (1992), 'Response', *Social Epistemology*, vol. 6 (1), pp. 45-55.
- Clifford, W.K., (1993) 'The ethics of belief', in L. Pojman (ed), *The Theory of Knowledge: Classic and Contemporary Readings*, Belmont, Wadsworth Inc., pp. 506-514.
- Foley, R. (1993), 'What am I to believe?', in S. Wagner and R. Warner (eds), *Naturalism: A Critical Appraisal*, Notre Dame, University of Notre Dame Press, pp. 147-163.
- Harman, G., (1994), 'Positive versus negative undermining in belief revision', in H. Kornblith (ed), *Naturalizing Epistemology*, second edition, Cambridge, MIT Press, pp. 317-337.
- Malcolm, N., (1993), 'Two types of knowledge', in L. Pojman (ed), *The Theory of Knowledge: Classic and Contemporary Readings*, Belmont, Wadsworth Inc., pp. 56-63.
- Polanyi M. (1958), *Personal Knowledge*, New York, Harper & Row.
- Polanyi, M. (1967), *The Tacit Dimension*, Garden City, Anchor Books.
- Stich, S. (1990), *The Fragmentation of Reason*, Cambridge, MIT Press.
- Taylor, S. and Brown, J. (1988), 'Illusion and well-being: A social psychological perspective on mental health', *Psychological Bulletin*, vol. 103 (2), pp. 193-210.
- Taylor, S., Collins, R., and Skokan, L., (1989), 'Maintaining positive illusions in the face of negative information: Getting the facts without letting them get to you', *Journal of Social and Clinical Psychology*, vol. 8 (2), pp. 114-129.
- Taylor, S. and Brown, J., (1994), 'Positive illusions and well-being revisited: Separating fact from fiction', *Psychological Bulletin*, vol. 116 (1), pp. 21-27.

⁷ I am indebted to Harold I. Brown for helpful comments on an earlier version of this paper.