

## Chapter 3. What is Knowledge?

LINDA ZAGZEBSKI

### Sections

- **1. Introduction: The Object of Knowledge and the Components of Knowledge**
- **2. Desiderata in Defining Knowledge**
- **3. The Traditional Definition of Knowledge and Gettier Objections**
- **4. A Definition of Knowledge**
- **5. Assessment of the Definition**

#### 1. Introduction: The Object of Knowledge and the Components of Knowledge

Knowledge is a highly valued state in which a person is in cognitive contact with reality. It is, therefore, a relation. On one side of the relation is a conscious subject, and on the other side is a portion of reality to which the knower is directly or indirectly related. While directness is a matter of degree, it is convenient to think of knowledge of things as a direct form of knowledge in comparison to which knowledge *about* things is indirect. The former has often been called *knowledge by acquaintance* since the subject is in experiential contact with the portion of reality known, whereas the latter is *propositional knowledge* since what the subject knows is a true proposition about the world. Knowing Roger is an example of knowledge by acquaintance, while knowing *that* Roger is a philosopher is an example of propositional knowledge.<sup>1</sup> Knowledge by acquaintance includes not only knowledge of persons and things, but also knowledge of my own mental states. In fact, the knower's own mental states are often thought to be the most directly knowable portion of reality.

Propositional knowledge has been much more exhaustively discussed than knowledge by acquaintance for at least two reasons. For one thing, the proposition is the form in which knowledge is communicated, so propositional knowledge can be transferred from one person to another, whereas knowledge by acquaintance cannot be, at least not in any straightforward way.<sup>2</sup> A related reason is the common assumption that reality has a propositional structure or, at least, that the proposition is the principal form in which reality becomes understandable to the human mind. So even though my experience of Roger leads me to know Roger, and my experience of my own emotions leads me to know what it is like to have such emotions, as a theorist I am hard put to answer the question "What is knowledge?" about either of them. The object of knowledge is more easily explained when it is a proposition. In this paper I will follow the usual procedure of concentrating on propositional knowledge, but in doing so I recognize that the theoretical convenience of propositional knowledge does not necessarily imply its greater importance.

Propositions are either true or false, but only true propositions link the knower with reality in the desired manner. So the object of knowledge in the sense of most interest to philosophers is usually taken to be a true proposition. The nature of truth, propositions, and reality are all metaphysical questions. For this reason epistemologists generally do not direct their major effort to these questions when writing as epistemologists, and so discussions of knowledge normally do not center on the object of knowledge, but rather on the properties of the state itself that make it a state of knowing. Accounts of knowledge, then, direct their attention to the knowing relation and focus more on the subject side of the relation than on the object side.

So far we have seen that knowledge is a relation between a conscious subject and some portion of reality, usually understood to be mediated through a true proposition, and the majority of epistemological attention has been devoted to the subject side of that relation. In the state of knowledge the knower is related to a true proposition. The most general way of characterizing the relation between the knower and the proposition known is that she takes it to be true, and this relation is standardly called the state of belief. The idea that the knowing state is a species of the belief state undergirds the almost universal practice in epistemology of defining knowledge as true belief plus something else. But this view can be disputed since the history of epistemic concepts shows that belief and knowledge were sometimes regarded as mutually exclusive epistemic states. This was either because it was thought that knowledge and belief have distinct objects, or because it was thought appropriate to restrict the range of belief to epistemic states evaluatively inferior to the state of knowledge.<sup>3</sup> The first worry has been settled to the satisfaction of almost all contemporary epistemologists by the adoption of the widespread view that propositions are the objects of belief as well as of knowledge and, in fact, the same proposition can be either known or believed. So a person may know today what he only believed yesterday – say, that his favorite team would win the game today. If this is right, there is no objection to the idea that knowledge is a form of belief on the grounds of a difference in their objects. The second worry can be settled by stipulating that to believe is to *think with assent*, a definition that comes from Augustine.<sup>4</sup> Since it is indisputable that to know propositionally is, among other things, to take a proposition to be true, and if to assent to a proposition just *is* to take it to be true, then on the Augustinian definition of belief it follows that knowing is a form of believing.<sup>5</sup>

It is reasonable, then, to maintain that knowledge is a form of belief, but this is not necessarily helpful to a quest for a definition of knowledge since the concept of belief is itself in need of definition, and there are some philosophers who maintain that the concept has outlived its usefulness.<sup>6</sup> Still, it is widely assumed that the concept of belief is clearer and less controversial than the concept of knowledge. And this has to be the case if the common practice of defining knowledge as a form of belief is to be not only true, but illuminating. I think the assumption is correct but I will not defend it.

From what has been said so far it follows that knowledge is a form of believing a true proposition. At this point in the process of defining knowledge it becomes much more difficult and more open to debate. All parties agree that knowledge is a good state, good at least in the sense of desirable, and perhaps also good in the sense of praiseworthy. But there are different kinds of praiseworthiness. Good looks, wit, and strength are desirable qualities and we praise others for having them, but we typically do not blame them when they lack such qualities. In contrast, we praise persons for having qualities like courage, kindness, or fairness, and we also blame them for their absence. This suggests that it is a requirement of the moral sense of the praiseworthy that it is a quality whose presence is praised and whose absence is blamed. But this is only roughly right since blame for absence is also missing at the high end of moral praiseworthiness. We praise persons for being noble or saintly, but we do not blame them when they are not.

Now it is indisputably true that knowledge is desirable, but is it praiseworthy, and if so, in what sense? Is its praiseworthiness closer to the praiseworthiness of good looks, the praiseworthiness of kindness, or the praiseworthiness of saintliness? It is significant that knowledge has not traditionally been treated as a moral concept, yet it has had many of the trappings of the moral – for example, the connection with epistemic duty and responsibility,

as when we criticize a person by saying she *ought* to know better, a criticism that is often accompanied by the type of distaste characteristic of the moral. Particular failings in knowledge are often attributed to qualities that have a decidedly moral tone, as when we say that a person is not fair to his intellectual opponents or is intellectually cowardly or is dogmatic. In each case the failing may be the explanation for the subject's lack of knowledge and he may be blamed for lacking knowledge because of this failing. Unfairness and cowardice are clearly qualities that have a moral sense, and dogmatism does also, although it is perhaps less obvious. A distinguishing feature of the dogmatic person is that he refuses to seriously entertain any evidence that might shake his belief; that is, nothing counts against it. But when we criticize a person for being dogmatic it is often very close to criticizing him for being a bigot. The response is akin to moral revulsion. In each of these cases, then, the failing is perceived to be like a moral one and if the subject lacks knowledge because of it, the lack of knowledge itself is perceived to be like a moral failing.

So sometimes the good of knowledge is treated like a moral good. A person is praised for its presence and blamed for its absence. But there are also instances of knowledge the lack of which is outside the realm of the blameworthy, and this indicates that moral concepts are not applicable. Obvious examples include perceptual and memory knowledge. It is usual these days to think that I know that I am looking at a yellow daffodil in ordinary circumstances in which I am looking at a yellow daffodil and form the belief that I am doing so, and all agree that that is a desirable state. It would be a stretch to say there is anything praiseworthy about it because it is so ordinary, and certainly the lack of perceptual knowledge in such circumstances due to a visual Abnormality is pitied rather than blamed. Of course, cases of knowledge by extraordinary perceptual acuity are praised and deserving of it, but the lack of perceptual knowledge in such cases is surely not blamed. So certain kinds of knowledge seem to be far removed from the moral realm.

One problem for the theorist is to reconcile these different senses in which knowledge can be good. Sometimes the good of knowledge is like natural goods, sometimes it is similar to moral goods, and sometimes it may even be thought to be like the noble. Major disputes over the definition of knowledge may turn on contrasting senses in which knowledge is good. According to the contemporary theory of reliabilism knowledge is true belief arising from a reliable truth-producing mechanism. This proposal makes the good of knowledge a natural good like that of beauty, wit, or strength. The traditional proposal that knowledge is true belief based upon good reasons is associated with the ethical concepts of responsibility, praise and blame. One is praised for believing the truth upon good reasons and blamed for not doing so. The idea that knowledge is noble comes from Plato.<sup>7</sup> In my judgment no definition of knowledge can succeed if it does not incorporate or at least adjudicate the senses of good used in these opposing types of theory.

## 2. Desiderata in Defining Knowledge

In section 1, I have given a general characterization of knowledge as a state of believing a true proposition in a good way. This much is widely accepted, although some of the deepest disputes over the definition of knowledge turn on the sense of good intended in this loose, preliminary definition. But more has to be settled before proceeding. The question "What is knowledge?" is not a question with a single clear purpose. To ask the question and to give an answer are human activities that arise out of a variety of human needs. If the question is a request for a definition, what sort of definition is wanted? In this section I want to address

this issue since some of the differences in accounts of knowledge arise from different aims in asking the question.

A definition can serve a number of different purposes, some practical, some theoretical. When we are defining knowledge one purpose might be the practical one of giving us directions for finding instances of knowledge in ourselves and in others, perhaps with the further aim of helping us to get it. A quite different purpose is the theoretical one of understanding where the *concept* of knowledge should be placed on a conceptual map that philosophers have already partially charted. This theoretical aim is intended to issue in a definition that is a necessary truth, whereas the practical aim may be satisfied by a contingent definition. Theoretical and practical purposes can sometimes be at odds.

A common theoretical purpose is to give what Locke called a *real definition*, a necessary truth that elucidates the nature of the kind of thing defined. Not all concepts defined by necessary truths have real definitions. For example, *bachelor* is defined as *an unmarried man*, and this is presumably a necessary truth, but no one thinks that bachelors constitute an independent kind of thing whose nature we want to investigate. In contrast, natural kinds like *human being*, *gold*, and *water* are thought to be good candidates for real definitions. In spite of the obvious dissimilarities between knowledge and these natural substances, it is common for philosophers to aim for a real definition of knowledge, although this is often not stated explicitly. I believe it is an aim that deserves attention, however, since it presupposes some disputable semantical and metaphysical views. Perhaps knowledge is not in an ontological category for which a real definition is possible. For example, no one would attempt a real definition of *rich*, *candy*, or *large plant*, and only some theorists would attempt a real definition of *food*, *intelligence*, or *virtue*. In some of these cases a contingent definition is probably sufficient, and it will be at least to some extent conventional. It is feasible to aim for a real definition of knowledge only if the concept of knowledge is not like the concept of a large plant. And even if it is closer to the concept of intelligence or of virtue, it is still undecided whether a real definition is attainable. In raising these questions about the purposes of definition my point is not to settle these matters, but to indicate how they prescribe what is wanted in an answer to the question "What is knowledge?"

The purpose of a definition might be reached by more than one method, so to criticize the method is not the same as to criticize the purpose.<sup>8</sup> The widespread purpose of giving a Lockian real definition issuing in a necessary truth can be attempted by more than one method. For decades the preferred method has been the method of truth condition analysis according to which putative necessary and sufficient conditions for being an instance of knowledge are proposed and tested by the method of counterexample (Chisholm, Klein, Plantinga, et al.). Recently this method has been under attack in general<sup>9</sup> and in the particular case of knowledge. For example, the recent theory of contextualism does not treat knowledge as a natural kind nor does it aim for a set of necessary and sufficient conditions.<sup>10</sup> Edward Craig seems to reject not only the method of truth condition analysis, but the purpose behind the method as well. Craig aims to identify the distinguishing features of knowledge by examining its pragmatic purpose in a community of informants. He seems to think of knowledge as closer to an artifact than a natural kind. He does not aim at a necessary truth and it does not bother him that not all knowers are good informants. The procedure of revising a definition by examining counterexamples is not part of his method.<sup>11</sup> Hilary Kornblith also rejects the method of truth condition analysis but not the purpose of giving a real definition. He sees knowledge as a natural kind, but he believes

empirical investigation can result in a necessary truth about it just as empirical investigation can lead us to discover necessary truths about physical or biological kinds like gold or water. Alvin Goldman uses truth condition analysis of the concept of knowledge, but he sees empirical methods as applicable to it since he believes concepts are psychological structures the contents of which are subject to empirical test. He doubts that the resulting definition will be a necessary truth and he does not aim for a real definition.<sup>12</sup> Goldman therefore retains the method but not the purpose.

Aristotle identifies a kind of definition that is “a formula exhibiting the cause of a thing's existence.” As an example he cites defining thunder by its efficient cause as what occurs “because fire is quenched in the clouds.”<sup>13</sup> Aristotle contrasts this with the kind of definition that purports to give the essential nature of a thing, its formal cause, and he suggests that the same thing can be defined both ways. It is interesting to compare Alvin Goldman's early causal theory of knowledge and more recent forms of reliabilism with the Aristotelian procedure of defining a thing through its causes.<sup>14</sup> Unlike Aristotle, these philosophers take their definitions to be rivals to definitions such as *justified true belief* that aim to elucidate the nature of the knowing state itself. Goldman suggests that sometimes the nature of a thing just *is* to be (efficiently) caused a certain way, e.g. sunburn, and he thinks of knowledge as like sunburn. But the essential nature or formal cause of most things is distinct from their efficient cause, so if knowledge is like most things, it could be defined either through its nature or through its efficient cause, and the two definitions would not compete. I do not know of anyone who has pursued this possibility.

There is much to be said for each of the theoretical and practical purposes mentioned above, but reflection might show us that some purposes do not make good philosophical sense. At any rate, conscious consideration of the purpose and method of definition in general can lead us to see alternatives that might otherwise slip by us when attempting to define knowledge. It is particularly desirable to question whether we should aim for a real definition since it is hard to determine whether knowledge is a single kind of thing for which a real definition is possible. Epistemologists almost always have the aim Plato had in the *Theaetetus*, where he says he is setting out to “bring the many sorts of knowledge under one definition” (148e). But do we know this aim is attainable?

The attempt to give a real definition of knowledge can be challenged by the fact that the concept of knowledge has been treated in many different ways in different periods of philosophical history. Is there really a single target of analysis about which all these accounts differ, or are some of them simply talking about different things? This question is particularly striking when we look at the differences in the rigor of the requirements for knowledge throughout philosophical history. According to some theories the conditions for knowledge are narrow and strict, whereas in others they are broad and loose. The philosophical tradition leans to the rigorist side, although the contemporary trend is in the opposing direction. It is now widely held that ordinary cases of perception and memory yield knowledge and that small children and possibly even animals have knowledge in these ways. But it is worth noticing how much this differs from a long line of rigorist accounts starting with that of Plato in the *Phaedo* and the *Republic*. Plato made knowledge a much loftier state than the ordinary, and the difference between his rigorist conception and the more lenient contemporary one may make us doubt that a real definition of knowledge is possible.

These same worries also arise when we examine the sense in which knowledge is good, addressed in section 1. We saw there that perceptual knowledge seems initially to be good

in a different sense from the good of the knowledge that requires reasons. Is it plausible to think that both phenomena are instances of the same kind of thing? Some philosophers have consciously divided the kinds of knowledge to reflect these differences.<sup>15</sup> The same problem arises with the treatment of skepticism. When the global skeptic says he does not know that he inhabits the planet Earth and I say he does know that he inhabits the planet earth, is it clear that we are disagreeing about something? Are we debating about the implications of a single concept or is there more than one concept that at times has gone by or has been translated by the term “knowledge”?<sup>16</sup> All of these worries may lead us to ask to what extent knowledge is a single phenomenon rather than a set of distinct phenomena, to what extent the boundaries of the kind are natural rather than set by convention, and to what extent “knowledge” is a term of philosophical art.

I believe we should begin by assuming that there is a single concept of knowledge about which philosophers have been debating for millennia and that we should aim for a necessary truth in our definition until forced to give up by continual failure in reaching the goal. I am less confident that knowledge is a single natural kind on a par with *water* or *gold*, but it is tempting to hope that that is the case. In any event, if knowledge is not a natural kind it is unlikely we will discover that unless we attempt to treat it as one. I will therefore tentatively accept the traditional aim of aspiring to a real definition of knowledge.

In addition to purposes and methods of definition, there are some common criteria for good definition that put limits on what will be acceptable, among them that a definition should not be *ad hoc*, that it should not be negative when it can be positive, that it should be brief, that it should not be circular, that it should utilize only concepts that are less obscure than the concept to be defined, and many others. I believe these criteria are good ones, although I will not examine them closely. Some of them serve the purpose of making a definition informative, a purpose that clearly goes beyond the aim of accuracy. The idea is that a definition is supposed to tell us something we didn't already know. We want a definition because of our failure to clearly grasp the concept to be defined. A circular definition does not do that since it uses the concept to be defined in the *definiens*, nor does one that uses other concepts as much in need of definition as the *definiendum*. Negative definitions fail in a more subtle way. They tell us what something is not, not what it is.<sup>17</sup> Of course, there are cases in which there is nothing more in a concept than the negation of another one. For instance, it is common to define *right act* as *an act that is not wrong*. In such cases we say that the negated concept in the *definiens* is conceptually more basic. I know of no reason to think that the concept of knowledge is like the concept of a right act, the negation of some other, more basic concept, so a negative definition of knowledge will probably be insufficiently informative.

The criterion that the concepts in the *definiens* should be less obscure than the *definiendum* may be one of the motives of those epistemologists who maintain that there should be no normative concept in the definition of knowledge. To be sure, they recognize that knowledge is good, but their aim is to define the sense in which knowledge is good in nonmoral and even nonnormative terms because of an assumption that nonnormative concepts are better understood than normative ones. This is often accompanied by worries about the ontological status of normative features of the universe. Normative facts and properties are thought to be puzzling in a way that descriptive facts and properties are not. I do not find this view plausible, but it is not a matter that can be settled without a deep investigation into the nature of normativity, and that would take us well outside the domain of epistemology. It is worth noting, however, that this criterion may

be at odds with one of the theoretical purposes of a definition already mentioned. If we want a definition to connect the concept to be defined with other key concepts in well-developed philosophical theories, the concepts that have a central place in such theories might turn out to be normative ones. And since knowledge is a normative concept, that is just what we would expect. If so, it might actually be an advantage to use these concepts in the definition. If ultimately it turns out that normative concepts are reducible to nonnormative ones, the demonstration that that is the case would be work for a further project.

One of the requirements for a good definition I have mentioned is that it not be *ad hoc*. This requirement is particularly telling when the method used is that of truth condition analysis. That method aims to make a definition counterexample-free, but the procedure of proposing a definition and then repeatedly repairing it in response to counterexamples can sometimes lead to a definition that is too obviously a response to problems in some *other* definition. This is one way a definition can fail to be either theoretically illuminating or practically useful.

In my view there is nothing wrong with there being a number of different definitions of knowledge of differing sorts, and it is helpful to keep their different purposes and methods in mind when one is compared with another. It is also a good idea to be sensitive to the difficulty in satisfying all the desiderata in a single definition. One may have to be sacrificed for the sake of another. For example, precision is clearly better than vagueness, but sometimes precision results in a definition of knowledge that is so long, cumbersome, and hard to remember that it serves neither the purpose of giving us theoretical understanding nor the practical one of giving us guidance in achieving it.

In the next section we will look at an important set of counterexamples that attacked a long-standing definition of knowledge. It will be helpful to keep in mind the various desiderata in a definition of knowledge while addressing that problem, since one moral that might be drawn from it is that there is a problem with the method of counterexample itself.

### 3. The Traditional Definition of Knowledge and Gettier Objections

So far we have concluded that knowledge is *good true belief*. Nobody would consider this an acceptable definition, however, because it adequately serves neither a theoretical nor a practical purpose. The concept of good is at least as much in need of analysis as the concept of knowledge. The definition does not specify what sense of good is intended, and even if it did, it does not provide us with the means to apply it to cases. On the other hand, it is brief, it is noncircular, and within the bounds of extreme vagueness, it is accurate.

Since believing is something a person does, beliefs have customarily been treated as analogous to acts, so beliefs are good in the sense in which acts are right. Right believing has traditionally been identified with justified believing. So knowledge is *justified true belief* (JTB).<sup>18</sup> Sometimes, but not always, this has been understood to mean *true belief for the right reasons*. For several decades the concept of justification has received an enormous amount of attention since it was assumed that the JTB definition of knowledge was more or less accurate and that the concept of justification was the weak link in the definition. For the most part these discussions proceeded under the assumption that the aim was to arrive at a necessary truth and that the method to be used in doing so was that of truth condition analysis. An important set of counterexamples to the JTB definition of knowledge were proposed by Edmund Gettier (1963) and led to many attempts at refining the definition

without questioning either the purpose or the method of definition. In this section we will look at the moral of Gettier's objection.

Gettier's examples are cases in which a belief is true and justified, but it is not an instance of knowledge because it is only by chance that the belief is true. Writers on Gettier normally do not say what they think is wrong with chance, but Aristotle does when he says, "To leave the greatest and noblest of things to chance would hardly be right."<sup>19</sup> Aristotle is here referring to *eudaimonia* or happiness, but his point is a general one about goods, at least great goods, and knowledge is surely a great good. It is incompatible with the value of knowledge that the aim of the knower, namely, getting the truth, occur by chance. This much has rarely been disputed even though, as we have seen, the sense of good intended has certainly been disputed.

In one standard Gettier example we are to imagine that Smith gives you plenty of evidence that he owns a Ford and you have no evidence against it. You then quite justifiably form the belief *Smith owns a Ford*. From that you infer its disjunction with *Brown is in Barcelona*, where Brown is an acquaintance whom you have no reason to believe is in Barcelona. Since the inference is justified, your belief *Smith owns a Ford or Brown is in Barcelona* is also justified. (Never mind what would possess you to form such a belief.) As it turns out, Smith is lying; he owns no Ford. But Brown is by chance in Barcelona. Your belief *Smith owns a Ford or Brown is in Barcelona* is therefore true and justified, but it is hardly something you know. Many examples of this kind have been proposed.<sup>20</sup>

As remarked above, it has often been noted that the problem in a Gettier case is that the truth is reached by chance; it is a kind of luck. But the structure of this case reveals that it is actually a case of double luck. It is mere bad luck that you are the unwitting victim of Smith's lies, and so it is only an accident that the kind of evidence that usually leads you to the truth instead leads you to form the false belief *Smith owns a Ford*. You end up with a true belief anyway because of a second accidental feature of the situation, a feature that has nothing to do with your cognitive activity. So an element of good luck cancels out the bad.

Some writers on Gettier have thought that the problem arises only for a restricted range of definitions, those according to which justified true belief means true belief based upon good reasons.<sup>21</sup> Since "justified" has sometimes meant "for the right reasons," this is understandable. Unfortunately, however, the problem is much more extensive than that. Given a couple of plausible assumptions already mentioned about what is required in an acceptable definition, it can be shown that Gettier problems arise for any definition in which knowledge is true belief plus something else that is closely connected with truth but does not entail it. It does not matter if the something else is a matter of believing for the right reasons or even if it is captured by the concept of justification. It need not even be anything accessible to the consciousness of the believer; for example, it may simply specify that the belief is produced by a reliable epistemic process or properly functioning faculties. All that is necessary is that there be a small gap between truth and the component of knowledge in addition to true belief in the definition. Call this component Q. In any such case a counterexample to the definition can be constructed according to the following recipe:

Start with a belief in the gap – that is, a belief that is false but is Q in as strong a sense of Q as is needed for knowledge. The falsity of the belief will not be due to any systematically describable element in the situation for if it were, such a feature could be used in the analysis of Q and then truth would be entailed by Q, contrary to the hypothesis. We may say that the falsity of the belief is due to some element of luck. Now amend the case by adding another element of luck, only this time one that makes the belief true after all. This second



element must be independent of Q so that Q is unchanged. We now have an example of a belief that is Q in a sense strong enough for knowledge and that is true, but that is not knowledge. The conclusion is that as long as the concept of knowledge closely connects the component Q and the component of truth but permits *some* degree of independence between them, no definition of knowledge as true belief plus Q will succeed.

A well-known attempt to avoid Gettier problems without giving up the essence of the JTB definition is to add defeasibility conditions to the definition. This idea was proposed when it was noticed that in typical Gettier cases the justified belief depends upon or otherwise “goes through” a false belief. In any event there is a fact unknown to the subject which would defeat her justification should she discover it. In our example it is the fact that Smith does not own a Ford. With this observation in mind, defeasibility theories add to the components of true belief and justifiedness the requirement that there are no truths, qualified in various ways, which when added to the reasons justifying the belief would make it no longer justified. In the strong defeasibility theory a belief is not knowledge unless there is *no* truth which when added to the reasons justifying the belief make it no longer justified. But that, of course, makes Q entail truth, so it is not a case in which there is a small gap between truth and the other conditions for knowledge. Weaker defeasibility theories do not close the gap between Q and truth, and they are still vulnerable to Gettier-style problems using the recipe I have proposed. That procedure allows us to produce counterexamples even when the belief does not depend upon a false belief and even when there is no false belief in the neighborhood.

The nature of induction allows us to produce examples of this kind. Suppose that Dr. Jones, a physician, has very good inductive evidence that her patient, Smith, is suffering from virus X. Smith exhibits all of the symptoms of this virus, and laboratory tests are consistent with the presence of virus X and no other known virus. Let us also suppose that all of the evidence upon which Jones bases her diagnosis is true, and there is no evidence accessible to her that counts against the diagnosis. The conclusion that Smith is suffering from virus X really is extremely probable on the evidence. But even the strongest inductive evidence does not entail the conclusion and so it is possible to make a mistake. Let us suppose that this is one of those cases. Smith is suffering from a distinct and unknown virus Y. Dr. Jones's belief that Smith is presently suffering from virus X is false, but it is justified and undefeated by any evidence accessible to her.

Now the recipe for generating a Gettier-style example tells us to add an additional feature of the situation that makes the belief true after all but without altering the other features of the situation. Let us say that besides suffering from virus Y, Smith has very recently contracted virus X, but so recently that he does not yet exhibit symptoms caused by X, nor is the laboratory evidence upon which Jones bases her diagnosis produced by X. So while the evidence upon which Dr. Jones bases her diagnosis does make it highly probable that Smith has X, the fact that Smith has X has nothing to do with that evidence. In this case Dr. Jones's belief that Smith has virus X is true, justified, and undefeated, but it is not knowledge.

This same example can be used to generate counterexamples to a host of other theories. Since even the strongest inductive inference can lead to a false belief, that false inductive belief will satisfy any requirement for the normative element of knowledge that is not necessarily connected to truth. But then we can always describe a situation that is identical except that the belief turns out to be true after all due to some extraneous aspect of the situation. In such a case the subject will not have knowledge but will satisfy the conditions of the definition.

We may conclude that the prevalent method of defining knowledge as true belief plus something else cannot withstand counterexample as long as there is a small degree of independence between truth and that something else. It follows that there must be a necessary connection between truth and the other conditions of knowledge in addition to truth, whatever they may be. In the first section we saw that these other conditions can be loosely defined as believing in a good way. So the sense in which knowledge is believing in a good way must entail truth.<sup>22</sup>

It must be observed, however, that the conclusion of this section is correct only if we accept some plausible assumptions mentioned in section II about desired features in a definition. That is because the problem can be avoided by giving a definition that is either *ad hoc* or is too vague to be useful. For example, the definition *knowledge is justified true belief that is not a Gettier case* is obviously not susceptible to Gettier-style examples, nor is the very general definition we started with: *Knowledge is good true belief*. The first definition is clearly *ad hoc* as well as negative, and we have already said that the second is not only much too vague, it uses a concept in the *definiens* that is at least as obscure as the concept of knowledge.

Since Gettier cases are those in which accidentality or luck is involved, it has often been suggested that knowledge is *nonaccidentally true belief*. This definition also is vague as well as negative and it has little practical import. It is not a counterexample to the point of this section, though, since nonaccidental truth entails truth. However, it has been shown by the Howard-Snyders that the component of knowledge in addition to true belief can be defined in a way that uses the concept of nonaccidentality but does not entail truth. Their definition is this: Knowledge is *true belief which is such that if it were true it would be nonaccidentally true*.<sup>23</sup> The idea behind this definition is the observation that the nonaccidental connection between the way in which knowledge is good and the truth need only obtain in the cases in which the belief is true since false beliefs are not candidates for knowledge. A false belief can have the property of being such that if it were true it would be nonaccidentally true and, hence, this property does not entail truth. This definition highlights an assumption I have made in my recipe for generating Gettier-style counterexamples, the assumption that if a false belief has the property that converts true belief into knowledge – the property Q – it is always possible that there be a belief that has Q but that is accidentally true. The Howard-Snyders' idea is to rule out that possibility in the definition of property Q. The resulting definition combines the defects of the previous ones. Like the definition *nonaccidentally true belief* it is vague, negative, lacks practical import, and has little to recommend it theoretically. Like the definition *justified true belief that is not a Gettier case* it is *ad hoc*. In addition, it has the problems that come with interpreting the truth conditions of the subjunctive conditional *If it were true it would be nonaccidentally true*. On the other hand, it at least appears to be nonnormative, a feature that ought to please those philosophers who aim for a definition of knowledge that contains no normative element.

Nonaccidentality is not a desirable element in a definition of knowledge, but it shows us something interesting about the process of defining knowledge. Nonaccidentality has been suggested as a component of knowledge not because it has been identified as a feature of paradigm cases of knowledge, but because *accidentality* is a feature of certain well-known cases of *nonknowledge*. The trouble is that the observation that an accidental connection between truth and component Q is insufficient for knowledge does not tell us what *is* sufficient for knowledge. Of course, the connection between truth and component

Q must be nonaccidental, but that is only the weakest thing we can say about it. Counterexamples are generally situations in which a defect in a definition is highlighted in an extreme form. But we should not conclude from that that anything less than the extreme defect is good enough. In this section we have seen that the connection between truth and the element of knowledge in addition to truth must be not only non-accidental, there must be no possibility at all of a gap between them. Closing the gap can be done in a variety of ways, not all of which require entailment, and I suggest that we should choose a way that respects the other desired features in a definition.<sup>24</sup> To avoid a definition that is *ad hoc* it is preferable that there be a conceptual connection between truth and the other element of knowledge. That is, knowledge is not only a good way of cognitively grasping the truth, it is also one in which the truth and the good way in which it is achieved are intrinsically related. That intrinsic relation ought to be explicit in the definition. Theories that have this feature have been proposed, although they have usually not recognized that the moral of Gettier demands it.<sup>25</sup>

In this section we have seen that if we accept some plausible requirements for an acceptable definition, Gettier cases arise whenever there is a gap between the truth and the other conditions for knowledge. This means that knowledge is not merely a summation of the component of truth and the other components. I have drawn the conclusion that we want a definition that makes a conceptual connection between truth and the sense in which knowledge is good. However, our analysis might support a more radical conclusion. The discussion of Gettier cases arises within the context of certain assumptions about the purpose and method of the definition. The aim is to get a necessary truth, perhaps also to get to a real definition, and the method used is that of truth condition analysis. But as we saw in section II, it is not obvious that these assumptions are warranted. In particular, the method of truth condition analysis can be and has been disputed. The Gettier problem might be interpreted as exhibiting the defects of such a method, thereby supporting the move to a different method altogether. As I have already stated, however, my own preference is the more conservative one of retaining the method of truth condition analysis but without letting the aim to make the definition counterexample-free dominate the list of desiderata adopted in section 2.

#### 4. A Definition of Knowledge

The conclusions of the first three sections of this paper give us a program for defining knowledge. Let us review them. In the first section I gave a rough definition of knowledge as believing a true proposition in a good way, and we saw that the sense of good intended in the concept of knowledge is a stumbling block to reaching a definition that encompasses both the cases of knowledge by perception or memory and the cases of knowledge that involve higher human abilities. The good of the former is similar to natural goods, whereas the latter are good in a sense that is close to the moral. The good of knowledge may sometimes even be like the most noble goods.

In section 2 I reviewed a number of different purposes and methods of defining knowledge and proposed that we try to satisfy as many of them as possible. But I will not try to satisfy the common purpose of eliminating all normative concepts from the definition. Since we know that the concept of knowledge is normative, it is a theoretical advantage if it can be related to central concepts in ethics since ethicists already have proposed theoretical structures in which these concepts have been analyzed. If it turns out that normative concepts are reducible to or supervene on nonnormative concepts, the demonstration that

that is the case would be an independent project. Meanwhile, one of my purposes will be to integrate the concept of knowledge into a background ethical theory.

In section 3 we looked at the moral of Gettier examples and concluded that the normative component of knowledge, the component that makes knowledge good, must entail the truth. Success in reaching the truth must be an intrinsic part of the sense in which each instance of knowledge is good.

I will propose a definition that attempts to meet all of these criteria. It should be clear from what has been said, however, that there is no unique way of doing so. In particular, the successful attainment of the theoretical purpose of locating the concept of knowledge on a background conceptual map depends upon what concepts are thought to be most theoretically salient, and that, in turn, depends upon which background theories have the most importance in the eyes of those asking the question, "What is knowledge?" But that, in turn, depends upon the resolution of deep issues in metaphilosophy. Should we try to embed the concept of knowledge in a background normative theory because it is a normative concept? Should we instead embed it in a background metaphysical theory on the assumption that metaphysics is more basic than epistemology? Or should we embed it in a scientific theory on the grounds that knowledge is a natural phenomenon? I have already said that I will take the first of these alternatives, but I have not argued for it and I can see many advantages in defining knowledge in terms of very different concepts from the one I have chosen. In fact, even if the purpose is to embed the concept of knowledge in a background ethical theory, the choice of theory will obviously depend upon one's position regarding the kind of ethical theory most likely to serve our theoretical and practical purposes.

The definition I will propose arises from a virtue theory of ethics. The complete theory includes intellectual as well as moral virtues within the same theory and aims to give a unified account of the morality of believing as well as of acting, but I will discuss only that part of the theory that underlies the normative concept I use in defining knowledge.<sup>26</sup> This is the concept of *an act of intellectual virtue*.

The concept of a virtue has a number of theoretical and practical advantages. Its proposed advantages in ethics are well known and I have argued elsewhere (1996) that there are parallel advantages in epistemology. In the last section we saw that the definition of knowledge must make success in reaching the truth an intrinsic aspect of that which makes knowledge good. The traditional concept of justification cannot serve this purpose, nor can any concept of a property of a belief. That is because no normative property of a belief guarantees its truth, at least no property the concept of which already has a history. But in Aristotle the concept of a virtue combines that of an admirable internal state with external success. At least, that is one way of interpreting Aristotle, and in any event, the concept of a virtue as used in ethics can be adapted to our need for a concept that makes an intrinsic relation between the good of a person's internal state – in this case, belief – and its success – in this case, the truth. So I suggest that it will be beneficial to move back a step from properties of beliefs to properties of persons in our search for a concept that attaches the good of knowledge to its truth.<sup>27</sup> Virtues are properties of persons. Intellectual virtues are properties of persons that aim at intellectual goods, most especially truth. Moral virtues are properties of persons that aim at distinctively moral goods such as the well-being of others. Since the concept of a virtue already has a rich history, if we can connect knowledge to virtue, that would be a theoretical advantage. In addition, the concept of a virtue has practical uses. Ordinary people speak of such individual virtues as kindness, fairness,

courage, open-mindedness, perseverance, generosity, discretion, and trust, and sometimes the same names are used for both moral and intellectual virtues. Furthermore, the evaluation of acts is often made in terms of the virtues or vices they express. The price of the practical usefulness of the concept of virtue and of the individual virtues may be a certain degree of conventionality in the application of the concept, although I will not discuss this aspect of the concept here. *Virtue* is not a technical concept, although it can be technically refined. I believe it is a virtue of the concept of virtue that it has both an extensive history in the philosophical literature and a wide use in ordinary discourse.

There are many accounts of the structure of a virtue. I will briefly summarize my own without argument.

A virtue has two components. The first is a motivational component and the second is a component of success in reaching the end of the motivational component. The motivational component of a virtue is a disposition to have an emotion that directs action towards an end. Each virtue has a distinctive motivational component with a distinctive end, but groups of virtues can be categorized by their ultimate ends. Most intellectual virtues have truth as their ultimate end.<sup>28</sup> Moral virtues have other ultimate ends. The success component of a virtue is a component of reliability in bringing about the end of the virtuous motivation. To take a few examples, the virtues of compassion, trust, and open-mindedness can be roughly defined as follows: The virtue of compassion is a trait that includes the emotion-disposition to alleviate the suffering of others and reliable success in doing so. The virtue of trust is the trait that includes the emotion-disposition to trust those and only those who are trustworthy, and reliable success in doing so. The virtue of open-mindedness is the trait that includes the emotion-disposition to be open to the views of others even when they conflict with one's own and reliable success in doing so. I suggest that the structure of all or, at least, most of the virtues can be defined by this pattern.

The concept of a virtue is important for character evaluation. When we say a person has a virtue we mean that she has a disposition to be motivated a certain way and to act a certain way in relevant circumstances, and in addition, is reliably successful in bringing about the end of her virtuous motive. But having a disposition to a motive does not mean she always has the motive in the relevant circumstances, and being reliably successful does not mean she is always successful. So the fact that she is virtuous does not entail that her individual acts and beliefs should be evaluated positively. At the same time, someone who is not virtuous may nonetheless be able to perform acts and have beliefs that are valuationally positive. The evaluation of acts and beliefs, then, requires further conditions.

Sometimes an act or belief has positive value simply because it is what a virtuous person would typically do in the circumstances, whether or not it is virtuously motivated. There is a sense of *right* in which we say a person has done the right thing in giving the correct change to a buyer even though he is not at all motivated by moral concerns. Similarly, there is a sense of *justified* in which we say a person has a justified belief in believing that the earth is a round even if he has not made the reasons for believing it his own. We also evaluate beliefs and acts from the aspect of the agent's motivation. An act or belief that is virtuously motivated deserves credit, although we almost always qualify it if it does not also involve doing/believing the right thing.

An act may be evaluated positively on both of these grounds and still not have everything we want morally in an act. So even when it is motivated properly and is what a virtuous person would do in the circumstances, it may fail in the aim of the act. When this happens the act lacks something morally desirable. Moral success is evaluated positively even though

that is to some extent out of the hands of the agent. It is one of the ways in which we are all victims of moral luck. So, for example, a person might be motivated by generosity and act in a way characteristic of generous persons in some particular circumstances, say by giving money to a beggar on the street, but if it turns out that the beggar is really rich and is playing the part of a beggar to win a bet, we would think that there is something morally lacking in the act. This is not, of course, to suggest that we would withhold praise of the agent, but her *act* would not merit the degree of praise due it if the beggar really were deserving. The same point applies to intellectual acts. A person may be motivated by intellectual virtues and act in a way intellectually virtuous persons act in attempting to get knowledge, but if she fails to get the truth, her epistemic state is lacking something praiseworthy. This means there is a kind of epistemic luck analogous to moral luck. As Thomas Nagel has remarked, the Nobel Prize is not given to people who are wrong.<sup>29</sup> Getting knowledge itself is a kind of prize, and it is in part the prize of being right.

In addition, mere success in reaching the end of the virtuous motive in the particular case is not sufficient for the highest praise of an act or belief even if it also has the other praiseworthy features just identified. It is important that success in reaching the end is *due* to the other praiseworthy features of the act. The end must be reached *because of* these other features. This is because there are ethical analogues to Gettier cases, although as far as I know, ethicists have not noticed this. Let me describe one such case.

Suppose a judge, weighing the evidence against an accused killer, determines by an impeccable procedure and motivated by justice that the man is guilty. We may assume that the judge not only does everything he ought to do, but he exhibits all the virtues appropriate in this situation. Nonetheless, even the most virtuous can make a mistake, just as we saw that even the most intellectually admirable can fail in an inductive conclusion in the case of Dr. Jones. Suppose this is one of those times. The accused is the wrong man. The fact that the judge makes a mistake is not due to any defect in him, whether moral or intellectual; it is simply bad luck. Obviously things have gone wrong, wrong enough that we would call the act a miscarriage of justice. The judge's act is not *an act of justice* even though we would not blame him for the error and would even praise him for *acting justly*. Nonetheless, the act itself is not deserving of the highest praise. It is lacking something morally important.

To get a Gettier-style problem we added an additional element of luck, a feature of good luck that cancels out the bad, and we can use the same procedure here. Suppose that the actual killer is secretly switched with the man the judge thinks he is sentencing so that the judge ends up accidentally sentencing the right man. One accident cancels out the other so that the end result is the desired one of punishing the culprit. In this situation I believe we would not give the judge's act the praise that would be due it if he had found the right man guilty in the first place. Of course we are relieved that the innocent man is not punished, but even though the end result is the one at which the judge was aiming and he was praiseworthy in both his motive and his actions, that is not sufficient to make his act the kind of act that deserves the highest moral praise.

The foregoing considerations show us that we need the concept of an act that gets everything right, an act that is good in every respect. And we have seen the elements that must be right or good in order to merit that evaluation. I call the concept that of *an act of virtue*. The definition is as follows:

An act is *an act of virtue A* if and only if it arises from the motivational component of *A*, is an act that persons with virtue *A* characteristically do in the circumstances, and is successful in bringing about the end of virtue *A* because of these features of the act.

The motivational component of *A* is a disposition. An act that arises from that disposition need not be consciously motivated by *A*, but it must be such that the explanation for the act would refer to it. An act that is characteristic of virtue *A* is an act that is not only what persons with virtue *A* would probably do in the circumstances, but it is an act that is a mark of the behavior of persons with that virtue.<sup>30</sup> The third component specifies that success in reaching the end must be because of the other two components. This needs further analysis. I know of no account of the *because of* relation that fully captures it, but I will have a bit more to say about it in the next section. It is important to notice that on this definition it is *not* necessary that the agent possess virtue *A* in order to perform an act of virtue *A*. One of Aristotle's conditions for virtue possession is that the trait must be deeply entrenched. If so, persons who are virtuous-in-training do not possess a given virtue, yet I see no reason to think they cannot perform acts of virtue, that is, acts that are as praiseworthy as an act can be with respect to the virtue in question.

There are acts of moral virtue and acts of intellectual virtue. We are concerned here with the latter. An act of intellectual virtue *A* is one that arises from the motivational component of an intellectual virtue *A*, is an act that persons with virtue *A* characteristically do in those circumstances, and is successful in reaching the truth because of these other features of the act.

The definition of knowledge I propose is as follows:

*Knowledge is belief arising out of acts of intellectual virtue*

At the beginning of this paper I mentioned that the common practice of concentrating on propositional knowledge in philosophical accounts of knowledge does not necessarily reflect its greater importance. All forms of knowledge involve contact of the mind with reality, however, and so a more comprehensive definition of knowledge that includes knowledge by acquaintance as well as propositional knowledge would be as follows:

*Knowledge is cognitive contact with reality arising out of acts of intellectual virtue*

Knowledge is generally not reached through a single act, but through a combination of acts of one's own, as well as through the acts of others and cooperating circumstances. We tend to think of knowledge as our own accomplishment, but this is rarely the case. The fact that our knowledge depends upon the knowledge and intellectual virtue of a host of other persons in our intellectual community, as well as a cooperating universe, makes it clear that we cannot expect to isolate the conditions for knowledge in some set of independent properties of the knower, much less a set of properties over which the knower has control. Epistemic luck permeates the human condition whether for good or for ill.

## 5. Assessment of the Definition

### A. Resolving the sense in which knowledge is good

At the end of the first section I said that no definition of knowledge can succeed unless it can resolve the different senses in which knowledge is good. Ordinary perceptual and memory knowledge seem to be good in a sense close to that of natural goods like beauty, wit, and strength. But sometimes knowledge is treated as a more elevated state, requiring

effort and skill. In these cases it seems to be good in a sense close to the moral. If so, it might turn out that knowledge is not a natural kind for which a real definition is possible. Perhaps investigation will reveal that there really are two distinct kinds of knowledge, just as investigation into the nature of jade revealed that what is called “jade” is really two distinct substances: jadeite and nephrite. I have not eliminated the possibility that ultimately this may happen, but I do not think we yet have a reason to bifurcate knowledge into two distinct kinds with separate analyses. The definition proposed in the last section can cover both kinds. In fact, I think it can even cover the highest sort of knowledge that is arguably in the realm of the noble.

An act of intellectual virtue has been defined within a background ethical theory in which *virtue* is the primary concept. “Virtue” is a term flexible enough to apply to more than moral traits, although the moral sense is no doubt the paradigm.<sup>31</sup> The definition of “an act of virtue” stretches the moral sense of the term in another way. An act of virtue is an act in which there is an imitation of the behavior of virtuous persons and success in reaching the end for that reason. More importantly for our present interest in interpreting the concept of an act of *intellectual* virtue, there is nothing in the definition that precludes that property from attaching to acts that are more or less automatic, as typically happens in perception and memory. The virtuous motivation from which an act of virtue arises need not be either conscious or strong, so ordinary epistemic motives will often be sufficient. In fact, nothing in the definition prevents the motivational component from applying to motives that are almost universal in some situations. The second component specifies that the act must be something that a person with the virtue in question would typically do insofar as he is expressing the virtue. But virtuous persons do not necessarily act in a way that is out of the ordinary, although, of course, they certainly do so in some circumstances. A person who has the virtue of attentiveness is as attentive as is necessary in situations of a given kind in order to reach the truth. A person who has the virtue of thoroughness examines the evidence as thoroughly as is necessary for the particular circumstances, and so on. Suppose a person with all the intellectual virtues is looking at a white wall in ordinary circumstances. Does she stare for a long time before forming the belief that there is a white wall in front of her? Does she undertake an investigation of the possibility that she is hallucinating or under the influence of drugs? Does she question trustworthy authorities on the subject of the color of walls? Of course not. To do so would exhibit a degree of intellectual scrupulosity tantamount to paranoia. But she *is* sensitive to any evidence that would lead her to suspect a defect in her perceptual ability or any peculiarities of the circumstances that would suggest a non-cooperating environment. Fortunately, most of the time she need not follow up on these possibilities. So to act like a person with intellectual virtue acts when judging the color of a wall is not a very difficult thing to do. And the same point applies to ordinary cases of belief based on clear memory. Typical true beliefs by perception or memory, then, satisfy my definition of knowledge. It is even possible that young children satisfy the definition as soon as they are old enough to know there is a difference between truth and falsehood and to be motivated to get the former.

So the definition can handle cases on the low end of knowledge. Its real advantage over other accounts, however, is at the upper end of knowledge. Stunning intellectual discoveries yield knowledge in a way that needs to be captured by any acceptable definition of the knowing state. Such knowledge is not merely the result of reliable processes or properly functioning faculties or epistemic procedures that have no flaw, as some epistemologists have suggested. They are the result of epistemic activities that go well beyond the



nondefective. They are, in fact, exceptionally laudatory. The concept of an intellectual virtue is well suited to the purpose of identifying knowledge in cases of this sort. A virtue is an admirable quality that goes beyond the minimum for being epistemically respectable. Some virtues go far enough beyond the minimum to reach the status of the highest goods. Creativity and originality of intellect are among those qualities associated with the high end of epistemic value, and an act of the virtue of originality is praiseworthy in the same way that acts of supreme generosity are praiseworthy. Such acts are truly exceptional. The definition of knowledge I have proposed, then, covers a range of cases from that of low-grade perceptual knowledge whose goodness is like natural goodness, to cases of beliefs based on evidence that are praised and blamed in the way we associate with the moral, to truly stellar intellectual achievements whose goodness is close to the noble.

### **B. How it escapes Gettier problems**

In section 3 I showed that unless we are willing to live with a very uninformative definition, Gettier problems result from any definition in which the sense in which knowledge is good does not entail truth. The concept of an act of intellectual virtue does entail truth, and so my definition is not guaranteed to fail in the way I have outlined for those theories susceptible to the double luck strategy. In the two cases we examined, that of the belief, *Smith owns a Ford or Brown is in Barcelona*, and the case of Dr. Jones and her diagnosis that Smith has virus X, the believer reaches the truth because of the feature of double luck I identified in these cases.<sup>32</sup> You and Dr. Jones do reach your respective beliefs because of your intellectually virtuous motivations and activities, but you do not reach the *truth* because of these features of the situation. This means that the concept of reaching *A because of B* is a key element of the definition. We all have intuitions about what it means for something to happen because of something else, but this concept is in need of further analysis and I do not know of one that is adequate. Some epistemologists have attempted counterfactual accounts of the component of knowledge in addition to true belief and, up to a point, whether the believer would arrive at the truth in close counterfactual circumstances can be used as a way of determining whether the truth is reached in the actual circumstances because of virtuous activity. So, for example, we might defend our claim that you do not get to the truth in the Ford and Barcelona case because of your virtuous motives and acts since in very similar circumstances you would have had the same motives and performed the same acts and failed to get to the truth. That would have been the case if Brown had not happened to be in Barcelona. Similarly, Dr. Jones would have reached a false belief in very similar circumstances even with her virtuous motives and acts. That would have happened if Smith had not happened to contract virus X just before she made her diagnosis. But looking at whether the believer reaches the truth in relevantly similar counterfactual circumstances is only a rough way of determining whether the truth is reached because of designated features of the act. It is certainly not a way of explaining what is *meant* by saying that the truth is reached because of these features. For example, there are no counterfactual circumstances in which a bachelor is not unmarried, but it would not be true to say that he is a bachelor *because* he is unmarried. The concept *A because of B* is not reducible to these counterfactual conditions. At best any such definition of *because of* will be a nominal definition.

### **C. Issues for further inquiry**

In the method of truth condition analysis the principal question is whether the definition is too broad (weak) or too narrow (strong). John Greco has objected to me that the definition

might be too weak in that it does not require the actual possession of intellectual virtue as a condition for knowledge. Since acts of intellectual virtue can be performed by agents whose virtuous behavior does not arise out of an entrenched habit, they cannot be trusted to act virtuously in similar circumstances in the future. Is it appropriate to attribute knowledge to them if they would not do the same thing in relevantly similar circumstances? I have said that to make the possession of the fully entrenched virtue a condition for knowledge is too strong since it rules out knowledge in children and unsophisticated adults, but Greco's point deserves further attention. It is likely that it would lead us into an investigation of the psychology of habit formation and the stability of the behavior of persons at early stages of acquiring intellectual traits. It also brings up the question of the extent to which we think an otherwise unreliable person can have knowledge because her behavior depends upon the reliability of other persons in her epistemic community.

It may also be objected that the definition is too strong. This is most likely to be raised against the motivational component of an act of intellectual virtue. Why think that the subject's motives have anything to do with whether she gets knowledge? This question highlights the differences between those who tend to think of knowledge as procedural and mechanical, and those who think of it as something for which we are responsible. My sympathies, of course, are with those in the latter category, but underlying the issue of whether responsibility extends to the cognitive sphere is a disagreement about the extent to which cognitive activity is voluntary. This suggests that deeper questions about human nature are at issue here.<sup>33</sup>

The definition as I have proposed it here meets many of the criteria for a good definition given in section 2, but it is vague and it clearly needs more extensive analysis. We have already seen the need for an account of the *because of* relation in the third component of the definition of an act of virtue. It also needs an account of motivation, as well as an account of acting in a way that is characteristic of a virtue, the first and second components of the definition. If an agent is doing what virtuous persons characteristically do in some circumstances, does that include having the relevant knowledge of the circumstances that virtuous persons have when they act? How far does that knowledge go? And if knowledge of circumstances is included in the account of the second component of an act of virtue, aren't we left with a circular definition since the concept of knowledge has been smuggled into the *definiens*?<sup>34</sup> There is also the matter of identifying and individuating the intellectual virtues. This is important not only because differing lists of the individual virtues and their analyses can result in accounts of knowledge that differ greatly in plausibility, but because it is possible that some of the virtues conflict. Virtue theories of ethics have this same problem. Aristotle's solution was to tie together the different virtues in the concept of *phronesis*, or practical wisdom, and I have attempted to use the same move with respect to the intellectual virtues.<sup>35</sup> But this move will not succeed unless it can be demonstrated that making every virtue relative to the judgment of a person with practical wisdom yields applications of the virtues to cases that are recognizably the same as the ones we have been using intuitively.

Many of these problems would have to be addressed by a detailed virtue theory of ethics anyway. There are, therefore, other motives in answering these questions besides the motive to define knowledge. A successful answer to them would serve a purpose in ethics as well as in epistemology. Other definitions of knowledge that meet the criteria I have described here would need to do the same thing, only they would refer to a different background theory in ethics, metaphysics, or cognitive psychology. The most detailed and

advanced of such theories will always have the advantage in providing a theoretical background for the definition of knowledge.<sup>36</sup>

1. Some philosophers have tried to reduce one of these forms of knowledge to the other.
2. See Kierkegaard's notion of indirect communication for his view on the way to communicate truth or subjectivity, which he believes is nonpropositional. This idea appears throughout his writings, but particularly in *Concluding Unscientific Postscript*.
3. Plato used both reasons for his view that the objects of knowledge (*episteme*) and belief (*doxa*) differ. See particularly the line analogy in the *Republic* 509d–511e, and the famous Allegory of the Cave 514a–518d.
4. Augustine, *Predestination of the Saints*, 5, trans. Dods, reprinted in Oats, *Basic Writings of St. Augustine*, 2 vols. (New York: Random House, 1948). The definition of believing as *thinking with assent* seems to make beliefs conscious occurrences and so to rule out belief in the dispositional sense, the sense in which we sometimes attribute beliefs to a person even when he is not thinking of them. But the Augustinian definition of believing can be extended to include a dispositional sense. Believing *p* dispositionally would be defined as having the disposition to assent to *p* when thinking of it.
5. But see H. A. Prichard for the view that knowledge and belief are mutually exclusive states and that we can tell the difference by introspection. Prichard says: "We must recognize that whenever we know something we either do, or at least can, by reflecting, directly know that we are knowing it, and that whenever we believe something, we similarly either do or can directly know that we are believing it and not knowing it" (*Knowledge and Perception*, Oxford: Clarendon Press, 1950, p. 86).
6. Stephen Stich, *The Fragmentation of Reason* (Cambridge, MA: MIT Press, 1990).
7. Plato calls knowledge the most important element in life (*Protagoras* 352d) and says that the only thing truly evil is to be deprived of it (345b).
8. For an interesting discussion of the purposes and methods of definition, see Richard Robinson, *Definition* (Oxford: Clarendon Press, 1950).
9. See John Pollock, "A Theory of Moral Reasoning," *Ethics* 96 (April 1986): 506–23. Pollock argues that concepts are not individuated by truth condition analysis but by what he calls their conceptual roles.
10. For examples of contextualism see David Annis, "A Contextualist Theory of Epistemic Justification," *American Philosophical Quarterly* 15 (1978): 213–19; Keith De Rose, "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52 (1992): 913–29; David Lewis, "Elusive Knowledge," *Australasian Journal of Philosophy* 74, no. 4 (December 1996): 549–67.
11. I am not sure that Craig sees himself as offering a definition of knowledge in his book. Nonetheless, he is attempting to answer the question, "What is knowledge?" So the striking difference between his purposes and methods and those that are more common in contemporary epistemology is directly related to the topic of this paper.
12. In personal correspondence.
13. *Posterior Analytics* II, 94a1–5.
14. Goldman 1967, 1986.
15. Ernest Sosa distinguishes between animal knowledge and reflective knowledge in "Intellectual Virtue in Perspective" and in "Reliabilism and Intellectual Virtue," in Sosa (1991).
16. William P. Alston has raised these same worries about the concept of justification in "Epistemic Desiderata," *Philosophy and Phenomenological Research* 53 (Sept. 1993): 527–51. But he says that he believes there is more commonality in the concept of knowledge (n. 15).
17. Plato uses this criterion for a good definition in the *Theaetetus* where Socrates says: "But the original aim of our discussion was to find out rather what it is than what it is not; at the same time we have made some progress, for we no longer seek for knowledge in perception at all" (187a).
18. Sometimes the JTB definition of knowledge has been compared with that of Plato in *Theaetetus* 201d in which Socrates considers and then rejects the proposal that knowledge is true belief (*doxa*) with an account (*logos*). It seems unlikely, though, that what Plato meant by a *logos* is very close to what contemporary philosophers mean by justification. In addition, Plato is not discussing propositional knowledge in that dialogue, but rather knowledge of persons or things.
19. *Nicomachean Ethics* 1109b25.
20. Bertrand Russell proposes an example of a stopped clock that is similar to Gettier cases in *Human Knowledge: Its Scope and Limits* (New York: Simon & Schuster, 1948), p. 154. But Russell uses it as a

counterexample to the proposal that knowledge is true belief. He does not seem to notice that if I have no reason to distrust my clock, my belief might be justified as well as true. This point was noticed by Israel Scheffler and is discussed by Robert Shope in *The Analysis of Knowing* (Princeton: Princeton University Press, 1983), pp. 19–20.

21. See Plantinga 1993: 36.
22. The argument of this section is taken from Zagzebski 1994 and Zagzebski 1996, part III, section 3.
23. Frances and Daniel Howard-Snyder, “The Gettier Problem and Infallibilism,” paper delivered at the Central Division meetings of the American Philosophical Association, May 1996. Sharon Ryan makes a similar proposal in “Does Warrant Entail Truth?” *Philosophy and Phenomenological Research* LVI (March 1996): 183–92, but she puts the conditional in the indicative mood. This is at least misleading since it suggests she intends a material conditional.
24. I have argued here that there must be a necessary connection between the component Q and truth. As I have stated my conclusion, however, Q must entail truth, although I have not argued that the connection must be as strong as entailment. Peter Klein has pointed out to me that a relationship of nomic necessity between Q and truth might be sufficient to avoid the double luck problem. That is, it might be sufficient that the gap between Q and truth is closed in every possible world with our causal laws. I will not pursue this approach here, however, since the relationship of entailment is the most straightforward way to make the required connection of necessity between the two components of knowledge, and I do not see any reason to think that theories vulnerable to the double luck formula I have outlined here would be any better off with a requirement of nomic necessity instead of entailment.
25. Three examples are Chisholm's early theory, Goldman's causal theory, and the strong defeasibility theory, already mentioned. Chisholm used the concept of making *p* evident in the definition of knowledge he proposes in the first edition of *Theory of Knowledge*, and he says there that whatever makes *p* evident must not also make evident a false proposition. This precludes the falsehood of *p*. Goldman's causal theory of knowledge had the truth condition built into the causal condition because he required that the subject does not know *p* unless the state of affairs *p* is appropriately causally connected to the belief *p*. This puts the truth of *p* in the causal condition. Since Goldman's later reliabilism does not build in the truth in this way, I assume that he was not motivated by the considerations I am giving here. In the strong defeasibility theory as expressed by Klein a belief is an instance of knowledge only if there is no true proposition which when added to the reasons that justify the belief makes the belief no longer justified. This condition entails the truth of the belief since if a belief *p* is false, *not p* is true, so there is a true proposition which if added to the subject's reasons for *p* entails the falsehood of *p*, namely, *not p*.
26. I have outlined a background ethical virtue theory in part II of *Virtues of the Mind*.
27. The move to properties of persons rather than properties of beliefs had already been made by reliabilists and earlier virtue epistemologists for different reasons. See my entry, “Virtue Epistemology,” in the forthcoming Routledge *International Encyclopedia of Philosophy* for a brief history of the development of virtue epistemology and its background in reliabilism.
28. There may be a few exceptions. Some virtues may aim at understanding rather than truth.
29. “Moral Luck,” in *Mortal Questions* (Cambridge: Cambridge University Press, 1979), n. 11.
30. In *Virtues of the Mind* I expressed the second component of the definition as follows: “it is something a person with virtue *A* would (probably) do in the circumstances” (p. 248). But what a virtuous person would probably do may not have anything to do with the virtue in question.
31. In *Virtues of the Mind* I argue that intellectual virtues are best treated as forms of moral virtue. The definition of knowledge does not depend upon this point, however.
32. Not all counterexamples in the Gettier literature have the double luck feature, although, of course, I have argued that cases with this feature can always be produced whenever there is a gap between truth and the other component of knowledge. But in every Gettier case there is some element of chance or luck.
33. I have discussed the issue of the voluntariness of cognitive activity in *Virtues of the Mind*, pp. 58–69.
34. This potential problem has been pointed out to me by Peter Klein.
35. *Virtues of the Mind*, part II, section 5.
36. I am grateful to Peter Klein, John Greco, and Richard Feldman for comments on an earlier draft of this paper, as well as to Hilary Kornblith and Alvin Goldman for correspondence while the paper was in progress.

### Cite this article

Zagzebski, Linda. "What is Knowledge?." *The Blackwell Guide to Epistemology*. Greco, John and Ernest Sosa (eds). Blackwell Publishing, 1998. Blackwell Reference Online. 25 May 2017 <[http://www.blackwellreference.com/subscriber/tocnode.html?id=g9780631202912\\_chunk\\_g97806312029126](http://www.blackwellreference.com/subscriber/tocnode.html?id=g9780631202912_chunk_g97806312029126)>