

“Knowledge as Credit for True Belief,” in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Michael DePaul and Linda Zagzebski, eds. (Oxford: Oxford University Press, 2004).

JOHN GRECO

KNOWLEDGE AS CREDIT FOR TRUE BELIEF

I begin by reviewing two kinds of problem for fallibilism about knowledge. The first is the lottery problem, or the problem of explaining why fallible evidence, though otherwise excellent, is not enough to know that one will lose the lottery. The second kind of problem is Gettier problems. I will argue that both kinds of problem can be resolved if we note an important illocutionary force of knowledge attributions: namely, that when we attribute knowledge to someone we mean to give the person credit for getting things right. Put another way, we imply that the person is responsible for getting things right. The key idea here is not that knowledge requires responsibility *in* one’s conduct, although that might also be the case, but that knowledge requires responsibility *for* true belief. Again, to say that someone knows is to say that his believing the truth can be credited to him. It is to say that the person got things right due to his own abilities, efforts and actions, rather than due to dumb luck, or blind chance, or something else.ⁱ

1. The lottery problem.

The lottery problem for fallibilism is stated nicely by Stewart Cohen.ⁱⁱ On the one hand, fallibilists want to say that there can be knowledge by inductive reasoning. Thus fallibilists define themselves against rationalists, who hold that only deductive grounds can give rise to knowledge. On the other hand, it seems that a ticket holder does not know that she will lose the lottery, even if the odds are heavily in favor of her losing. So here is the problem for fallibilism: how is it that in general one can know through inductive grounds, but in the lottery case one fails to know, though one’s inductive grounds are excellent?

To sharpen the problem, consider two cases of inductive reasoning.

Case 1. On the way to the elevator S drops a trash bag down the garbage chute of her apartment building. A few minutes later, reasoning on the basis of past experience and relevant background knowledge, S forms the true belief that the bag is in the basement garbage room. Of course her grounds for so believing are merely inductive: it is possible that the trash bag somehow gets hung up in the chute, although this is extremely

unlikely.ⁱⁱⁱ

Case 2. S buys a ticket for a lottery in which the chances of winning are ten million to one. A few minutes later, reasoning on the basis of past experience and relevant background knowledge, S forms the true belief that she will lose the lottery. Of course her grounds for so believing are merely inductive: it is possible that she buys the winning ticket, although this is extremely unlikely.

Many will have the intuition that S knows in Case 1 but not in Case 2. But how so, given that her reasons are excellent in both cases? This is what the fallibilist needs to explain.

a. Nozick's tracking account.

The difficulty of the problem is illustrated by some accounts that fail to solve it. According to Robert Nozick, S knows p just in case S's believing p tracks the truth.^{iv} Complications aside, this amounts to the following:

- S knows p just in case
1. p is true
 2. if p were true then S would believe p
 3. if p were false then S would not believe p.

Nozick's account does a good job explaining why S does not know in the lottery case: if S were going to win the lottery, she would still believe that she is going to lose. However, the tracking account rules incorrectly in the garbage chute case. This is because it is possible, although an extremely unlikely occurrence, that the trash bag gets hung up in the chute. But if the bag did get hung up, S would still believe that the bag is in the basement, and so S fails to satisfy clause (3) of Nozick's account. In fact, the garbage chute case and one's like it were originally formulated as counterexamples to Nozick's account.^v

b. Sosa's safety account.

Let us say that one's belief is sensitive to the truth just in case it satisfies Nozick's clause (3). Ernest Sosa suggests that a belief is better safe than sensitive, where a belief is safe just in case one would have it only if it were true.^{vi} Complications aside, Sosa's suggestion amounts to the following:

- S knows p just in case
1. p is true
 2. S believes p

3. S would believe p only if p were true. Alternatively, S would not believe p unless p were true.

Sosa's clause (3) is ambiguous between the following two interpretations:

Strong Safety: In close worlds, always if S believes p then p is true. Alternatively, in close worlds never does S believe p and p is false.

Weak Safety: In close worlds, usually if S believes p then p is true. Alternatively, in close worlds, almost never does S believe p and p is false.

But now there is a problem. If Sosa means to endorse strong safety as a condition on knowledge, then his account does no better than Nozick's with Case 1: S would believe that the bag is in the basement even if it were hung up in the chute, and so S fails to satisfy a strong safety condition on knowledge.^{vii} This suggests that Sosa means his safety condition to be interpreted as weak safety. But now his account rules incorrectly that there is knowledge in Case 2: It is true that in close worlds, usually if S believes she will lose the lottery then it is true that she will lose, and so in the lottery case S satisfies a weak safety condition.

That is how Sosa's account rules on the lottery case when complications are aside. When complications are not aside, Sosa adds that there are other conditions required for knowledge, or at least for full-blooded, reflective knowledge. For example, in one place Sosa says that a belief's safety "must be fundamentally through the exercise of an intellectual virtue," where an intellectual virtue is a reliable or trustworthy source of truth.^{viii} In another place he says: "For reflective knowledge one not only must believe out of virtue. One must also be aware of doing so"^{ix} But it seems clear that these added requirements do not make a difference in the lottery case, for there S *does* believe through reliable inductive reasoning, and might even be aware that she does.

c. Cohen's contextualism.

Finally, consider Cohen's own solution to the lottery problem. According to Cohen, the problem is solved by recognizing that attributions of knowledge are sensitive to context, and, more specifically, that the standards for knowledge are sensitive to context. We have knowledge in cases of ordinary

inductive reasoning, such as that employed in the garbage chute case, because the standards that are operative in ordinary contexts are low enough to admit such cases as counting for knowledge. We do not have knowledge in the lottery case, however, because in that context the standards for knowledge are raised—the possibility of winning the lottery becomes salient, and our inductive evidence, as good as it is, does not rule out this possibility.^x

I do not wish to deny Cohen's general point that the standards for knowledge are sensitive to context-- it seems to me that they are. What is less clear is how this is supposed to solve the lottery problem for fallibilism. The problem is that Cohen gives us no explanation why the standards for knowledge should get raised so high in the lottery case. More specifically: he gives no explanation why the standards should be raised beyond S's capacities to meet them. Cohen is quite explicit that he means to remain within the framework of fallibilism. Moreover, in the lottery case it is stipulated that S has excellent (although fallible) reasons for believing that she will lose. So why, on a fallibilist account of knowledge, does S fail to know that she will lose? To be clear, I am not claiming that S *does* know in the lottery case—I agree that she does not. My complaint is that nothing in Cohen's account explains *why* S does not know.

The same problem can be viewed from a different angle. Cohen says that when S reasons about the odds, the very form of her reasoning makes the possibility that S wins salient. And once made salient, Cohen says, that possibility cannot be ruled out. But again, why can't it be? Why isn't S's reasoning about the odds good enough to rule out the possibility of winning, even once made salient? It has been stipulated that S has excellent reasons for thinking she will not win the lottery, so why doesn't she know that she will not win? In sum, Cohen's contextualism does not explain what it was supposed to explain: given that we are fallibilists about knowledge, and given that we think inductive grounds are good enough to know in other cases, why are S's grounds not good enough to know in the lottery case?^{xi}

2. Gettier problems.

It has long been understood that fallibilist accounts of justification give rise to Gettier problems. For

example, consider the following.

Case 3. On the basis of excellent (although fallible) reasons, S believes that her co-worker Mr. Nogot owns a Ford: Nogot testifies that he owns a Ford, and this is confirmed by S's own relevant observations. From this S infers that someone in her office owns a Ford. As it turns out, S's evidence is misleading and Nogot does not in fact own a Ford. However, another person in S's office, Mr. Havit, does own a Ford, although S has no reason for believing this.^{xii}

Clearly S does not know that someone in her office owns a Ford. The problem for the fallibilist is to explain why this is so, given that S has excellent evidence for this true proposition.

3. A proposal for resolving the two problems for fallibilism.

We may distinguish two questions one might try to answer when giving an account of knowledge. The first is the "What is knowledge?" question. This question asks what conditions a person must satisfy to count as knowing. The second is the "What are we doing?" question. This question asks what illocutionary act is being performed when we say that someone knows. I will have more to say about the "What is knowledge?" question below. But I think that the key to solving our two problems for fallibilism lies in the "What are we doing?" question.

So what are we doing when we attribute knowledge to someone? Clearly, we might be doing any number of things. But one of the central functions of knowledge attributions is to give credit for true belief. When we say that S knows p, we imply that it is not just an accident that S believes the truth with respect to p. On the contrary, we mean to say that S gets things right with respect to p because S has reasoned in an appropriate way, or perceived things accurately, or remembered things well, etc. We mean to say that getting it right can be put down to S's own abilities, rather than to dumb luck, or blind chance, or something else. But then this gives us a resource for solving the two problems for fallibilism reviewed above. For in the lottery case, it *does* seem to be just a matter of chance that S gets it right when S believes that she will lose the lottery. And in Case 3 (the Gettier case), it seems just a matter of luck that S gets it right that someone in her office owns a Ford. In the garbage chute case, however, we think that it is due to S's good reasoning that she gets

things right—we give S credit for arriving at the truth in this case, and we are therefore willing to say that she knows.

This is the main idea that I want to pursue.^{xiii} The idea needs to be developed, however. For one, my treatment of the two problems remains largely intuitive so far. More importantly, the main idea I am proposing is faced with the following problem. First, I have said that we are willing to give S credit for her true belief in cases of knowledge, but not in the lottery case or in Gettier cases. Second, I have said that this is because in the former, but not in the latter, S's arriving at the truth is due to her own efforts and actions, or more exactly, to her own abilities. But it is not clear why the various cases can be distinguished in this way. Consider that in the lottery case, S uses excellent reasons to draw the conclusion that she will lose. And in the Gettier case, S reasons flawlessly to the conclusion that someone in her office owns a Ford. But then it seems that S arrives at the truth because of her abilities in all of the cases above, and not only in the cases where we judge that S has knowledge. Clearly, more needs to be said.

By way of saying more, I will draw on some important work in moral theory. Specifically, I want to look at Joel Feinberg's fascinating discussion of attributions of moral blame.^{xiv} From this it will be easy enough to construct a general theory of credit attribution, and a theory of intellectual credit attribution in particular. With this groundwork in place, it will be possible to explain why it is appropriate to give S credit for her true belief in cases of knowledge, and appropriate to withhold such credit in the lottery case and Gettier cases. More specifically, it will be possible to explain why, in cases of knowledge, it is appropriate to say that S gets things right *because* of her own abilities, whereas in the lottery case and Gettier cases, it is appropriate to deny this.

a. Feinberg on blaming.

Feinberg's account of moral blaming takes off from the following central idea: When we attribute blame to a person for some occurrence, part of what we are doing is assigning causal responsibility to that person for the occurrence. Put another way, when we blame S for X's occurring, we imply that S figures importantly into a correct causal explanation of why X occurred. To get further

insight into our practices of blaming, therefore, Feinberg makes some observations about the pragmatics of causal explanations in general. One important aspect of causal explanation language is this: In general, when we say that Y occurs because X occurs, or that Y's occurring is due to X's occurring, we mark out X's occurring as a particularly important or salient part of the causal story behind Y's occurring. For example, to say that the fire occurred because of the explosion is not to say that the explosion caused the fire all by itself. Rather, it is to say that the explosion is a particularly important part, perhaps the most important part, of the whole story. Or to change the example: to say that the fire occurred because of S's negligence is not to say that S's negligence caused the fire all by itself. Rather, it is to say that S's negligence is a particularly salient part, perhaps the most salient part, of the set of relevant factors that caused the fire.^{xv}

What determines salience? Any number of things might, but Feinberg cites two kinds of consideration as particularly important. First, among the various necessary parts of a complete causal process, an explanation will often pick out what is abnormal in the case, or what is contrary to expectations. For example, we will say that sparks caused the fire if the presence of sparks in the area is not normal. That explanation would misfire, however, if we were trying to explain the cause of a fire in a welding shop, where sparks are flying all the time. Or suppose that a white elephant walks into a room and causes a panic. Of course the white elephant entering the room is not sufficient all by itself to cause the panic—it would not if the room were part of a zoo and the people inside were animal trainers. But if the room is a place where white elephants are not expected to be, and if the people inside *are* as we would expect them to be, we have no trouble picking out the elephant as “the” cause of the commotion.

Another major factor governing salience is our interests and purposes. For example, often when we are looking for something's cause we are looking for something that we can manipulate to good effect. If the thing to be explained is smoke coming from the engine, for example, we will look for a part that needs to be replaced. Here it is perfectly appropriate to say that the cause of the smoke is the malfunctioning carburetor, although clearly a faulty carburetor cannot cause smoke all by itself. Or witness the various explanations of New York City's plunging crime rate. The police attribute it

to good policing, the Mayor attributes it to broader social policy, and opposing politicians attribute it to things over which the Mayor has no control, such as an upturn in the national economy. Of course any honest person would admit that the crime rate is influenced by all of these things. But different people have different interests, and so highlight different parts, of the causal conditions that were together sufficient for the drop in crime.

None of this need be insincere, nor is salience something that is merely subjective. The argument over “the” cause of the drop in crime is an argument over which of many causal factors should be deemed most important, given our collective interests and given what we know about human behavior, how we want to spend limited resources, what policies have what costs, etc. Likewise, a correct explanation of what caused a panic focuses on what is in fact abnormal about the situation. As was noted above, if the presence of elephants in the room were not in fact abnormal, then a correct explanation would have to focus elsewhere. In sum, correct causal explanations pick out some salient necessary part of a total set of causal factors, where salience is determined by a number of factors, including what is in fact normal and what are our actual interests.^{xvi}

We may now revisit Feinberg’s point that attributions of blame imply causal explanations. This is most obvious in cases where the thing for which the person is blamed is a consequence of the person’s actions. For example, to blame someone for the fire is to imply that her actions caused it: she is the one who struck the match, or who did not pay attention, or who did pay the arsonist. Alternatively, we can blame a person for the action itself, implying that she herself was the action’s cause, or perhaps that her choice was or her efforts were. As Feinberg notes, the distinction between blaming someone for her action and blaming someone for a consequence of her action is often merely verbal. For example, we can say either “She caused the fire by striking the match,” or “She started the fire.” Likewise, “She caused his death by poisoning his food,” substitutes for “She killed him.”

Finally, Feinberg argues that when we blame someone for an action we imply that the action reveals something important about the person himself: “In general, I should think, a person’s faulty act is registerable only if it reveals what sort of person he is in some respect about which others have

a practical interest in being informed.”^{xvii} Feinberg’s position is perhaps too strong on this point; it would seem that people can be rightfully blamed for actions that are out of character. Nevertheless, there does seem to be a *kind* of blame that Feinberg is right about. In other words, even if not all blaming implies that the person’s action reveals a faulty character, there is a strong sort of blame, which is common enough, that does. Moreover, this strong sort of blame has a counterpart in a strong sort of credit. Often enough, credit for an action implies a judgement about the person as well, implying not only that the person is responsible for doing something good, but that this is a manifestation of virtuous character.

Putting all this together, Feinberg’s account of blame for an action can be summed up as follows.

- A person S is morally to blame for action A only if
- a. A is a morally faulty action,
 - b. A can be ascribed to S, and
 - c. A reveals S’s faulty moral character.

Feinberg concludes that attributions of blame share the same pragmatics as causal explanations. His argument for this emphasizes clause (b) of the above account: attributing blame involves ascribing action, and ascribing action involves causal citation. What I want to emphasize, however, is that clause (c) acts the same way. Clause (c) also insures that attributions of blame involve causal citation, for what does it mean to say that an action reveals character, other than that the action results from character? In other words, clause (c) can be read:

- c’. S did A *because* S has a faulty moral character.

This might seem too strong, and it is if we read (c’) as saying that S’s character was sufficient all by itself to cause S’s action. Similarly, it is too strong if we read (c’) as saying that, given S’s character, S had to do A. But it is not too strong if we remember the pragmatics of causal explanation language reviewed above. For according to that account, to say that S’s action is a result of her character is to say that S’s character is an important part, perhaps the most important part, of

the story. Taken this way, (c') is not too strong at all, but rather reflects our common sense attitudes about the sources of human action. The fact is, we cite character in explanations of human behavior all the time, as when we say that he made the remark because he is insensitive (as opposed to having a bad day), or that she failed to spend the money because she is cheap (as opposed to hard up for cash at the moment).^{xviii} Feinberg's analysis reveals that such explanations are implied in attributions of blame, or at least in attributions of a certain sort of blame. And this implies that attributions of blame (of that special sort) will inherit the pragmatics of causal explanations. Clearly, any action will be the result of a number of factors, including a person's character. But sometimes we want to say that character is particularly salient—that it is an important part, perhaps the most important part, of the story behind why the person acted as he did.

b. A general theory of credit attribution.

Feinberg's account of moral blaming can easily be broadened in two ways. First, I have already noted that the counterpart of blame for an action is credit for an action. In fact, we can use credit as the general term, and talk about positive credit (i.e. praise) and negative credit (i.e. blame) for an action. Second, there are kinds of credit other than moral.^{xix} For example, we credit athletes for athletic feats and thinkers for intellectual ones. Accordingly, I propose the following as a general theory of credit.

A person S deserves credit of kind K for action A only if

- a. A has value of kind K,
- b. A can be ascribed to S, and
- c. A reveals S's K-relevant character. Alternatively: S's K-relevant character is an important necessary part of the total set of causal factors that give rise to S's doing A.

Two examples will illustrate this account.

Case 4. Ken Griffey Jr. runs full speed toward the center field wall, leaps with outstretched glove, and catches the ball while diving to the ground. The home team crowd, just robbed of a game winning double, shakes their respective heads in admiration of Griffey's

spectacular catch.

Case 5. Griffey Jr. runs full speed toward the center field wall, trips, and falls face down on the ground. The ball bounces off his head, goes straight in the air, and comes down in his glove. The home team crowd, just robbed of a game winning double, shakes their respective heads in disgust.

In both cases, the action in question has clear athletic value—catching the ball before it hits the ground is essential to winning baseball games. Moreover, in both cases the catch is ascribable to Griffey—we can be sure that a broadcaster announcing the game will be yelling, “Griffey caught the ball! Griffey caught the ball!” But only in Case 4 will Griffey be given credit for catching the ball, and that is because in Case 4 Griffey’s catching the ball is the result of his relevant character; i.e. his great athletic abilities. In Case 5 Griffey’s catching the ball was just dumb luck, and so the home team crowd is not just a bunch of sore losers. They are right to be disgusted.

A similar phenomenon occurs when a poor fielder makes a spectacular catch. In this case he will be given credit of a sort—he will get pats on the back from his teammates and applause from the crowd. But it won’t be the same kind of credit that Griffey gets. *Griffey* makes spectacular catches all the time—*his* catches manifest his great skills. Not so when Albert Belle makes such a catch. If the catch is difficult, it is almost just good luck that he makes it. And opposing fans will treat it that way, withholding the credit they would readily give to Griffey. Or consider Bucky Dent’s infamous home run to knock the Red Sox out of the play-offs. To this day Boston fans do not give Dent credit for the home run or the Yankees credit for the win. Dent was just a singles hitter, and his fly ball would have been a routine out in any park but Fenway. The home run was just bad luck, Boston fans think, having little to do with Dent’s abilities as a hitter.

Finally, it is interesting that the case can be viewed in more than one way. Yankee fans *do* give Dent credit for the home run—to them, he will always be considered a great hero. This is because Yankee fans don’t think that Dent’s home run was just a matter of luck. Their thinking emphasizes the idea that *some* luck is always involved in sports, but Dent got his bat on the ball, he was strong enough to muscle it out there, he was able to take advantage of Fenway’s short left field. In other words, their account of the home run downplays luck and emphasizes Dent’s abilities. This is a

common enough phenomenon in sports: Losers try to deny credit by emphasizing the role of luck, and winners try to take credit by putting the emphasis back on ability. Hence the attractive but dubious claims that “Good teams make their own luck” and “It all comes out even in the end.”

c. A theory of intellectual credit attribution.

When we attribute knowledge to someone we imply that it is to his credit that he got things right. It is not because the person is lucky that he believes the truth—it is because of his own cognitive abilities. He figured it out, or remembered it correctly, or perceived that it was so. Applying the account of credit attribution above, we have:

- S deserves intellectual credit for believing the truth regarding p only if
- a. believing the truth regarding p has intellectual value,
 - b. believing the truth regarding p can be ascribed to S, and
 - c. believing the truth regarding p reveals S’s reliable cognitive character.
Alternatively: S’s reliable cognitive character is an important necessary part of the total set of causal factors that give rise to S’s believing the truth regarding p.

And hence:

- S knows p only if believing the truth regarding p reveals S’s reliable cognitive character.
Alternatively: only if S’s reliable cognitive character is an important necessary part of the total set of causal factors that give rise to S’s believing the truth regarding p.

We are now in a position to apply these results to the two problems for fallibilism.

4. The lottery problem solved: Chance undermines credit.

The application to the lottery problem is initially straightforward: knowledge attributions imply attributions of intellectual credit for true belief, and intellectual credit implies that the true belief is the result of S’s own intellectual abilities. But here as in other cases, salient chance undermines credit. In the lottery case, but not in the garbage chute case, it seems just a matter of chance that S believes the truth. In the garbage chute case, but not in the lottery case, S’s true belief is appropriately credited to her, i.e. to her intellectual abilities.

This was our initial treatment of the case. But this initial treatment gave rise to the following problem: S employs admirable inductive reasoning no less in the lottery case than in the garbage

chute case. In both cases, therefore, S's abilities make up a necessary part, but only a part, of the whole story regarding S's believing the truth. The present account of credit solves this problem. For it is only in the garbage chute case that S's abilities are a *salient* part of the story. In the lottery case, what is most salient is the element of chance.

Why does the element of chance become salient in the lottery case? I would suggest that the very idea of a lottery has the idea of chance built right into it. Here is the way we think of the lottery case: First, S reasons on the basis of excellent grounds that she will lose the lottery. Second, the lottery is held and reality either does or does not match up with S's belief—it's just a matter of chance. Notice that things are different if S believes that she lost the lottery because she reads the results in the newspaper.^{xx} Here again her evidence is merely inductive, but now the role of chance does not play an important part in the story. Here is the most natural way to think of the newspaper case: First the lottery is held and the facts are fixed. Second, S infers from a reliable source that she has lost the lottery. Now it is not just a matter of chance that she believes the truth—she believes the truth because she has the good sense to believe what she reads about the results.

Two more cases help to explore the implications of the present account.

Case 6. On the basis of excellent evidence, S believes that her friend will meet her in New York City for lunch. However, S's friend has bought a ticket in the lottery, and if he were to win then he would be in New Jersey to collect his winnings. In fact, the friend loses the lottery and meets S for lunch.^{xxi}

Intuitively, people can know things about their friends even when, unknown to them, their friends have bought tickets in the lottery and would behave differently if they won. On the other hand, there is some intuitive pull toward thinking that S does not know in Case 6, since if her friend were to win the lottery then he would not meet her for lunch. The present account explains these conflicting intuitions by distinguishing contexts of attribution. In contexts where considerations about the lottery are not important, the salience of chance is low and so attributions of knowledge can still be appropriate. However, in contexts where considerations about the lottery are important, the salience of chance rises and therefore knowledge attributions are undermined. Handling Case 6

this way allows us to say that knowledge is closed under known entailment. Taking a page from the book of standards contextualists, we may note that there is no single context relative to which a) S knows that her friend will meet her for lunch, b) S knows that her friend will meet her for lunch only if he loses the lottery, and c) S does not know that her friend will lose the lottery. If the context does not make chance salient, then relative to that context S can know that her friend will meet her for lunch. But if the context does make chance salient, then relative to that context S knows neither that her friend will meet her for lunch nor that he will lose the lottery.^{xxii}

Now consider another case.

Case 7. S is visiting from another culture and knows nothing about lotteries. However, over a long period of time S observes many people exchanging dollar bills for small slips of paper, which they invariably discard after a brief examination. On the basis of excellent inductive reasoning, S concludes that the next person in line will soon discard his slip of paper. And in fact S's belief is true.^{xxiii}

This case threatens a counterintuitive result: that relative to S's context, S knows that the next person in line will discard his ticket. Put another way, if S were to express a knowledge claim, then, relative to S's own context of attribution, his claim would be true. This result threatens because S understands nothing about the workings of lotteries. And therefore, it seems, chance could not be a salient factor in S's context. However, this way of interpreting the case assumes that salience is a psychological phenomenon—that salience is a function of where someone's attention is, or how he is thinking about things. This is admittedly one meaning of the term, but it is not the one that we want here. On the contrary, I have argued that salience is a function of a) what is in fact normal or abnormal in the case, and b) what interests and purposes are in fact operative in the context of attribution. Therefore, whether something is salient relative to a context of attribution, in the sense intended here, is not a question about how anyone is thinking.

What we are looking for is not psychological salience but explanatory salience. The way to identify what is salient in this sense is to ask where a correct explanation would have to focus. As noted above, a correct explanation of a car fire might properly focus on sparks coming from the engine, whereas a correct explanation of a fire in a welding shop must focus somewhere else. But

this has nothing to do with how people are thinking about the cases. Rather, it has to do with what is in fact normal and abnormal in cars and welding shops. Likewise, if a drunk driver runs a light and is involved in a collision, then almost always a correct explanation will focus on the driver's impaired condition. Given the interests and purposes that are usually in place, other explanations of the collision are almost always inappropriate.

These considerations explain why it is correct to say that S does not have knowledge in Case 7, even relative to S's own context of attribution. For although the role of chance has no psychological salience for S, it does have explanatory salience; that is, any correct explanation of why S believes the truth would have to refer to the lottery. This is because playing the lottery is pretty much the only thing going on here. Relative to almost any context we can imagine, including S's own, an explanation that did not refer to the lottery would leave the most important thing out. This is in contrast to Case 6. In that case any number of factors contribute to S's believing the truth regarding her friend meeting her for lunch. And relative to most contexts that we can imagine, the fact that her friend loses the lottery will play no important part in a correct explanation of why she believes the truth.

Finally, it will be helpful to compare the present account with standards contextualism. First, both accounts are contextualist; that is, they both make the truth conditions of knowledge claims relative to the context of attribution. But the way this works in the two accounts is different. According to standards contextualism, the context of attribution determines the standards for knowledge, so that standards are higher or lower relative to different contexts. On the present account, however, the context of attribution determines the salience of various contributing causal factors, thus determining responsibility for true belief. Standards are not raised or lowered according context; rather, responsibility for a complex event (someone's believing the truth) is creditable or not creditable to the believer according to context.

Second, I argued above that familiar versions of standards contextualism do not explain why S's inductive reasoning does not allow her to know that she will lose the lottery. The standards contextualist says that in the lottery case the standards for knowing get raised because the possibility

of winning becomes salient. But this does not tell us why S's reasoning fails to meet those standards, even if raised. The present account does better in this respect. The very idea of a lottery involves the idea of chance, and so we have an explanation why chance is salient in cases where the lottery is salient. We can then apply a familiar general principle of credit attribution: that salient chance undermines credit.

5. Gettier problems solved: abnormality trumps interest.

Recently I have argued that the following conditions are necessary for knowledge.^{xxiv}

S knows p only if

1. S's believing p is *subjectively* justified is the following sense: S's believing p is the result of dispositions that S manifests when S is trying to believe the truth, and
2. S's believing p is *objectively* justified is the following sense: The dispositions that result in S's believing p make S reliable in believing p. Alternatively, the dispositions that result in S's believing p constitute intellectual abilities, or powers, or virtues.

The thinking behind the subjective justification condition is that knowledge requires that belief be appropriate from the knower's point of view. More specifically, the knower must have some awareness that a belief so formed has a good likelihood of being true. Some authors have required that the knower *believe* that this is so, but I have resisted this way of understanding the kind of awareness in question. It seems that people rarely have beliefs about the genesis of their beliefs, and so it would be too strong to require that they always have one in cases of knowledge. Accordingly, I have stated the subjective justification condition in terms of dispositions to believe rather than actual beliefs, or even dispositional beliefs. People manifest highly specific, finely tuned dispositions to form beliefs in some ways rather than others. And this fact, I take it, amounts to an implicit awareness of the reliability of those dispositions.

The objective justification condition can be interpreted along the lines of Sosa's weak safety condition above: to say that S is reliable in believing p is to say that, at least in relevantly close worlds, usually if S believes p then p is true. Notice that clause (2) makes the knower the seat of reliability: it says that S is reliable, not that some process is, or that S's method is, or that S's evidence

is. The thinking here is that knowledge must be grounded in the knower. That is, it must be grounded in the knower's own abilities, rather than in a process or method that might be engaged in accidentally, or on evidence that might be trusted on a whim.

I now want to add a third condition on knowledge, understanding knowledge attributions to imply attributions of intellectual credit, and understanding intellectual credit along the lines above. I propose that adding this third condition makes the three sufficient as well as necessary for knowledge.

3. S believes the truth regarding p *because* S is reliable in believing p. Alternatively: The intellectual abilities (i.e. powers or virtues) that result in S's believing the truth regarding p are an important necessary part of the total set of causal factors that give rise to S's believing the truth regarding p.

I claimed above that clause (3) allows us to handle a range of Gettier cases. Intuitively, in cases of knowledge S's reliable character has salience in an explanation of how S comes to get things right. In Gettier cases, S's reliable character loses its salience in favor of something else. What we want now is an account of why this is so. Hence we have two questions before us: First, why is it that S's cognitive abilities have salience in cases of knowledge. Second, why is it that they do not have salience in Gettier cases.^{xxv}

To answer the first question we may turn to a point that has been emphasized by Sosa. He writes,

All kinds of justification involve the cognitive or intellectual virtues, faculties, or aptitudes of the subject. We care about justification because it tends to indicate a state of the subject that is important and of interest to his community, a state of great interest and importance to an information-sharing social species. What sort of state? Presumably, the state of being a dependable source of information over a certain field in certain circumstances.^{xxvi}

For present purposes Sosa's point may be glossed this way: For beings like us, dependability (or reliability) of cognitive character is of fundamental social importance. But this implies that cognitive character will have a kind of default salience in the explanation of true belief. Unless something else is present that trumps the salience of S's cognitive abilities, those will be the most important part of

the story, or at least a very important part of the story.

We may now turn to the second question: Why is it that S's cognitive abilities are not salient in Gettier cases? Gettier cases seem to fall into three categories in this respect. First, there are cases where S's abilities are not salient in an explanation of true belief because S's abilities are not involved at all. Here is one such case.

Case 8. Charlie has excellent reasons to believe that he will make his connection in Chicago. However, Charlie does not recognize these reasons for what they are, nor does he base his belief on them. Rather, he believes that he will make his connection because he is overcome with wishful thinking at the prospects of seeing his fiancée. As it turns out, his belief is true.

Here it is straightforward that Charlie does not believe the truth because of his intellectual abilities, since his abilities are not involved at all in the production or maintenance of his belief. As I have described the case, he believes the truth entirely out of wishful thinking.

Second, there are cases where S's belief is produced by faculties that would normally be considered abilities, but where S is not reliable in the environment where those faculties are used. Here are two.

Case 9. Henry is driving in the countryside and sees a barn ahead in clear view. On this basis he believes that the object he sees is a barn. Unknown to Henry, however, the area is dotted with barn facades that are indistinguishable from real barns from the road. However, Henry happens to be looking at the one real barn in the area.^{xxvii}

Case 10. Rene is the victim of an evil demon who causes him to be systematically deceived. Despite admirable intellectual care and responsibility, almost all of Rene's beliefs are false. However, his belief that he is sitting by the fire, which is based on appropriate sensations and background beliefs, happens to be true.^{xxviii}

In Case 9 S's belief is the result of perception, and normally S's perception would constitute a cognitive virtue, i.e. a reliable ability or power. However, reliability is relative to an environment, and S's perception is not reliable relative to the environment in the example. Similar things can be said about Case 10—Rene's beliefs are caused by his cognitive faculties and in a normal

environment these faculties would constitute abilities or powers. But in Rene's actual environment his faculties are unreliable.

Finally, there are Gettier cases where S does use reliable abilities or powers to arrive at her belief, but where this is not the most salient aspect of the case. Case 3 regarding the Ford is one. Here are two others.

Case 11. A man *takes* there to be a sheep in the field and does so under conditions which are such that, when a man *does* thus take there to be a sheep in the field, then it is *evident* to him that there is a sheep in the field. The man, however, has mistaken a dog for a sheep and so what he sees is not a sheep at all. Nevertheless, unsuspected by the man, there *is* a sheep in another part of the field.^{xxx}

Case 12. Smith looks at his watch and forms the belief that it is 3 o'clock. In fact, Smith's watch is broken. But by a remarkable coincidence, it has stopped running exactly twelve hours earlier, and so shows the correct time when Smith looks at it.^{xxx}

In all three of these cases, reliable cognitive character gives rise to the belief in question – Lehrer's office worker reasons from her evidence, Chisholm's sheep watcher trusts his vision, and Russell's time teller trusts his watch. But in none of these cases does the person believe the truth because of reliable character. On the contrary, an explanation of how the person comes to get things right would have to focus somewhere else. In Case 3, S believes the truth because someone else in her office owns a Ford. Chisholm's sheep watcher believes the truth because there is a sheep in another part of the field. Russell's time teller believes the truth because his watch has stopped running exactly twelve hours before he looked at it. In each of the cases, there is something else that is more salient than S's cognitive abilities in the explanation of how S came to get things right.

Why is that so in this third category of Gettier cases? I propose that, in at least many such cases, there is something odd or unexpected about the way that S comes to believe the truth, and that the salience of the abnormality trumps the default salience of S's cognitive abilities. We have already seen that the salience of abnormalities is a general phenomenon. When a white elephant enters the room, all bets are off for other explanations regarding why the room emptied out. The odd or the unexpected undermines the salience of other factors involved in the event. In fact, Gettier cases are

often told so as to increase this effect. Typically, there is a first part of the case where an entirely normal scene is described. Then comes the “big surprise” part of the case, where we get the odd twist: what S sees is actually a dog, but there happens to be a sheep in another part of the field; S’s evidence is misleading, but someone else in the office owns a Ford.

Not all abnormalities undermine the salience of cognitive character, however. For example, an unlikely coincidence reminds the detective of evidence he has neglected, and this missing piece of the puzzle allows him to solve the crime. Less dramatically, an unusual noise causes someone to turn her head, and something else comes into plain view. We would like an account of which abnormalities undermine the salience of character in the explanation of true belief and which do not. Here is a suggestion that seems to handle the above cases well: abnormalities in the way one *gets* one’s evidence do not undermine credit, whereas abnormalities regarding the way one gets a true belief, given that one has the evidence that one does, do undermine credit. Put another way, in cases where something unusual does take away the salience of character, it seems just a matter of good luck that S ends up with a true belief, *even given that she has the evidence that she does.*^{xxxii}

We may understand this point in a way that generalizes to other kinds of credit: we may say that, in general, situational luck does not undermine credit. In other words, luck in the way that one gets into one’s situation does not undermine credit for what one does once in that situation. This, I take it, is one of the lessons of the literature on moral luck. A general prohibition on luck shrinks the sphere of moral responsibility to nothing, and hence an adequate theory of moral responsibility must allow for the influence of some kinds of luck but not others.^{xxxiii}

In sum, the present account goes some way toward explicating the content and pragmatics of knowledge attributions, although we must recognize that it does not go all the way. It goes some way because it analyzes knowledge in terms of credit attribution and credit attribution in terms of the content and pragmatics of causal explanations. It does not go all the way, however, because the pragmatics of causal explanations, and especially those concerning the salience of partial causes, are not fully understood. For example, the account given above in this regard is clearly only a partial account.^{xxxiiii} Nevertheless, our intuitions about knowledge seem to follow our intuitions about

causal explanation: In cases of knowledge, we think that S believes the truth because she is reliable, whereas in Gettier cases we think that S's believing the truth must be explained in some other way.

6. Concluding remarks.

I would like to end with two considerations that lend further support to the account of knowledge that has been offered here. First, the account predicts certain phenomena that do in fact take place. Second, the account suggests a solution to the value problem for knowledge.

a. Conflicting knowledge attributions.

If the preceding account of knowledge is correct, then there should be arguments over knowledge attribution that mirror arguments over other kinds of credit attribution. More specifically, there should be cases where disagreements over whether someone has knowledge reflect different emphases on the relative importance of S's contribution to arriving at the truth. This is in fact the case. For example, consider the following conversation between a gambler and his wife.

Gambler: I told you so, Honey! I knew that horse would win!

Wife: You knew no such thing, you idiot! That horse could have lost-- you nearly threw away our rent money again!

Gambler: No—I knew she would win! I've been watching that horse for months now, and I knew that she would run great in these conditions.

Wife: You son-of-a-bitch.

The conversation lends itself to the following interpretation: The gambler is trying to take credit for his true belief that the horse would win, and his wife is trying to deny him credit. In doing so, he tries to emphasize his ability to pick winning horses, and she tries to emphasize the role of luck.

In the context set by the above example I would tend to agree with the wife. But perhaps there are other contexts relative to which the gambler's claim to know is correct. Consider a conversation that might take place among gambling buddies.

First Gambler: How did you know that horse was going to win?

Second Gambler: Are you kidding? I've seen her run in mud before, and the only horse in

the race that could touch her was pulled just before post time. When I saw that, I ran to the betting window.

First Gambler: You son-of-a-bitch.

In this context the claim to know is accepted by all involved. Perhaps this is because gamblers are deluded about the relative importance of luck and ability in picking horses. Alternatively, in the present context the role of luck is taken for granted, and so all emphasis is rightly put on the relative abilities among the gamblers. If the “amateur” in the next row also picked the winner, he *wouldn't* be given credit for his true belief.

b. The value problem.

Recently, Linda Zagzebski has called attention to the value problem for knowledge.^{xxxiv} An adequate account of knowledge, she points out, ought to explain why knowledge is more valuable than mere true belief. In closing, I will suggest how the present account solves that problem.

In Book II of the *Nicomachean Ethics*, Aristotle makes a distinction between a) morally virtuous action and b) action that is merely as if morally virtuous. One important difference, says Aristotle, is that morally virtuous actions “proceed from a firm and unchangeable character.” [II.4] Moreover, it is morally virtuous action, as opposed to action that is as if virtuous, that is both intrinsically valuable and constitutive of the good life: “human good turns out to be activity of soul exhibiting excellence.” [I.7] The same point holds for intellectually virtuous action, where the distinction between “virtuous action” and “action as if virtuous” translates to a distinction between knowledge and mere true belief. Following Aristotle, therefore, we get an answer to the value problem: As is the case regarding moral goods, getting the truth as a result of one’s virtues is more valuable than getting it on the cheap.^{xxxv xxxvi}

ⁱ The idea that knowledge entails credit for true belief can be found in Ernest Sosa, “Beyond Skepticism, to the Best of our Knowledge,” *Mind* 97 (1988): 153-89, and *Knowledge in Perspective* (Cambridge: Cambridge University Press, 1991); and in Linda Zagzebski, *Virtues of the Mind* (New York: Cambridge University Press, 1996), and “What is Knowledge?” in John Greco and Ernest Sosa, eds. *The Blackwell Guide to Epistemology* (Oxford: Blackwell Publishers, 1999). More explicitly, Wayne Riggs argues that in cases of knowledge “we deserve credit for arriving at true belief non-accidentally.” See his “Reliability and the Value of Knowledge,” *Philosophy and Phenomenological Research*, forthcoming.

ⁱⁱ Stewart Cohen, “How to Be a Fallibilist,” *Philosophical Perspectives* 2 (1988): 91-123, and “Contextualist Solutions to Epistemic Problems: Skepticism, Gettier and the Lottery,” *Australasian Journal of Philosophy*, forthcoming.

ⁱⁱⁱ The example is from Ernest Sosa, “Skepticism and Contextualism,” *Philosophical Issues* 10 (2000): 1-18. We can imagine that the bag’s getting hung up is extremely unlikely because everything would have to go just right for that to occur, including the trajectory of the bag, its contents, the distribution of its weight, etc.

^{iv} See Robert Nozick, *Philosophical Explanations* (Cambridge, MA: Harvard University Press, 1981).

^v Here is another counterexample to Nozick’s account, due to Jonathan Vogel: “Suppose two policemen confront a mugger, who is standing some distance away with a drawn gun. One of the officers, a rookie, attempts to disarm the mugger by shooting a bullet down the barrel of the mugger’s gun. (I assume that the chances of doing this are virtually nil.) Imagine that the rookie’s veteran partner knows what the rookie is trying to do. The veteran sees him fire, but is screened from seeing the result. Aware that his partner is trying something that is all but impossible, the veteran thinks (correctly as it turns out) [that the] rookie missed.” From “Tracking, Closure, and Deductive Knowledge,” in Steven Luper-Foy, ed., *The Possibility of Knowledge* (Totowa, New Jersey: Rowman and Littlefield, 1987).

^{vi} See Ernest Sosa, “Skepticism and Contextualism”; “How Must Knowledge be Modally Related to What is Known,” *Philosophical Topics* 26 (1999): 373-384; and “How to Defeat Opposition to Moore,” *Philosophical Perspectives* 13 (1999): 141-155.

^{vii} Here I am assuming that there is a close world in which the bag gets hung up in the chute. If that seems wrong, we can invoke Vogel’s rookie cop example from note 5. There it seems uncontroversial that there is a close world where the rookie’s bullet enters the mugger’s barrel.

^{viii} “How Must Knowledge be Modally Related to What is Known?”, p. 383, note 7.

^{ix} *Knowledge in Perspective*, p. 278.

^x “How to Be a Fallibilist,” pp. 106-7.

^{xi} My claim is not that standards contextualism cannot explain the lottery case in principle. Rather, I restrict myself to the weaker claim that Cohen’s contextualism does not in fact explain it. What are the prospects for other versions of standards contextualism? The trick, of course, is for the standards contextualist to explain why S does not have knowledge in the lottery case, while at the same time preserving the intuition that S does have knowledge in other cases of inductive reasoning. But this will be hard to do. For example, Keith DeRose argues that S has knowledge if her belief matches the truth out to the nearest world where a salient alternative possibility is actual. However, the matching requirement insures that DeRose’s account rules incorrectly in the garbage chute case and in the rookie cop case from note 5. This is because these cases are designed so that not-p worlds are very, very close, and so no matter how weak the

standards for knowledge are being set, S's belief will not match the truth far enough out into alternative possible worlds. In other words, no matter how close the nearest world where a salient possibility is actual, S's belief will not match the truth out to that world. See Keith DeRose, "Solving the Skeptical Problem," *Philosophical Review* 104 (1995): 1-52.

^{xii} The example is from Keith Lehrer, "Knowledge, Truth and Evidence," *Analysis* 25 (1965): 168-75.

^{xiii} I first suggested a solution to Gettier problems along these lines in my review of Jonathan Kvanvig's *The Intellectual Virtues and the Life of the Mind*, in *Philosophy and Phenomenological Research*, vol LIV, no 4 (December 1994), pp. 973-976. Linda Zagzebski develops the idea in a different direction in *Virtues of the Mind* and in "What is Knowledge?," and Keith Lehrer develops a similar idea in *Theory of Knowledge*, 2nd edition (Boulder, CO: Westview Press, 2000). Earlier than any of this, Sosa writes that in cases of knowledge one's belief must "non-accidentally reflect the truth of P through the exercise of . . . a virtue." However, he does not suggest that this idea can be used to address Gettier problems. See his "Beyond Skepticism, to the Best of our Knowledge," p. 184. See also *Knowledge in Perspective*, p. 277.

^{xiv} Feinberg's discussion takes place over three papers, all of which are collected in *Doing and Deserving: Essays in the Theory of Responsibility* (Princeton: Princeton University Press, 1970). The papers are "Problematic Responsibility in Law and Morals," "Action and Responsibility," and "Causing Voluntary Actions." Page numbers that follow correspond to *Doing and Deserving*. My account of Feinberg's account of blaming is a reconstruction—I have taken parts of what he says from each of his three papers and put them together in a way that suits my present purposes.

^{xv} It is tempting to follow Feinberg and to put things this way: When we say that Y occurs because X occurs, or that Y's occurring is due to X's occurring, we mark out X's occurring as a particularly important or salient part of a sufficient condition for Y's occurring. (177) This assumes, however, that all causes can be understood as sufficient conditions. Since I do not want to deny either a) the possibility of agent causation or b) the possibility of indeterminate causation, I employ the looser language above.

^{xvi} Another example: Sports fans will argue endlessly over why we lost the big game. Was it because we gave up too many points or because we didn't score enough? Obviously, the outcome of a game is a function of both points allowed and points scored. The real argument here is over what was the most important factor in the loss. And *that* is a function of what one can normally expect, what could have been done differently, etc.

^{xvii} Feinberg, p. 126.

^{xviii} Some recent work in social psychology suggests that common sense is flawed in this respect. For example, see Lee Ross and Richard Nisbett, *The Person and the Situation* (New York: McGraw-Hill, 1991). For a persuasive argument against such a conclusion, see Michael DePaul, "Character Traits, Virtues and Vices: Are there None?," *Proceedings of the Twentieth World*

Congress of Philosophy, Volume IX: Philosophy of Mind and Philosophy of Psychology, forthcoming.

^{xix} Feinberg's own discussion is at times aimed at other kinds of blame.

^{xx} This point is made by Cohen in "How to Be a Fallibilist."

^{xxi} This kind of case is discussed by Cohen in "How to Be a Fallibilist," and by Gilbert Harman in *Thought* (Princeton: Princeton University Press, 1974).

^{xxii} Cohen and DeRose have both argued that contextualists need not run into closure problems. See Cohen, "How to Be a Fallibilist" and DeRose, "Solving the Skeptical Problem." For an early discussion of relevant issues, see Gail Stine, "Skepticism, Relevant Alternatives, and Deductive Closure," *Philosophical Studies* 29 (1976) 249-261.

^{xxiii} This sort of case was raised by Phillip Quinn in discussion.

^{xxiv} See my *Putting Skeptics in Their Place* (New York: Cambridge University Press, 2000); and "Agent Reliabilism," *Philosophical Perspectives* 13 (1999): 273-296.

^{xxv} What is a Gettier case? Zagzebski has argued that all Gettier cases are ones where bad epistemic luck is cancelled out by good epistemic luck. For example, in Case 3 S's evidence is deceptive (bad luck), but someone else in S's office owns a Ford (good luck). This analysis suggests that we can treat Gettier cases in the same way that we treated the lottery problem above—we can say that Gettier cases involve salient luck, and salient luck undermines credit. In my opinion this assessment is correct, but it is not as informative as we would like. This is because in Gettier cases, to say that S believes the truth because of good luck is very close to saying that S believes the truth for some reason other than her own abilities. And although that seems true, it is not more informative than what we already have—which is that clause (3) is violated. (In the lottery case, we said that the clause is violated due to the role of salient chance. That explanation *is* informative, however, because we have an independent grasp of what we mean by chance in a lottery.)

^{xxvi} *Knowledge in Perspective*, pp. 281-2.

^{xxvii} The example is from Alvin Goldman, "Discrimination and Perceptual Knowledge," *Journal of Philosophy* 73 (1976): 771-791. Reprinted in Goldman, Alvin, *Liaisons: Philosophy Meets the Cognitive and Social Sciences* (Cambridge, MA: MIT Press, 1992).

^{xxviii} The example is from Sosa, "Perspectives in Virtue Epistemology: A Response to Dancy and BonJour," *Philosophical Studies* 78 (1995): 221-235.

^{xxix} The example is quoted from Roderick Chisholm, *Theory of Knowledge*, 2nd edition (Englewood Cliffs, NJ: Prentice-Hall, Inc., 1977), p. 105.

^{xxx} The watch example is due to Bertrand Russell in *Human Knowledge: Its Scope and Limits* (New York: Simon and Schuster, 1948). The example is cited by Chisholm in *Theory of Knowledge*, where he points out that it is a Gettier case.

^{xxxi} The distinction employed here is similar to Mylan Engel's distinction between veridical and evidential luck, in Mylan Engel, Jr., "Is Epistemic Luck Compatible with Knowledge?," *The Southern Journal of Philosophy* XXX, 2 (1992): 59-75. I thank Michael Bergmann for pointing out to me that Engel's distinction is helpful in the present context.

^{xxxii} In this regard, see Thomas Nagel, "Moral Luck" in *Mortal Questions* (Cambridge: Cambridge University Press, 1979); Margaret Walker, "Moral Luck and the Virtues of Impure Agency," *Metaphilosophy* 22 (1991): 14-27; and John Greco, "A Second Paradox Concerning Responsibility and Luck," *Metaphilosophy* 26 (1995): 81-96.

^{xxxiii} We can imagine Gettier-type cases, for example, where the problem is not abnormality in the way that S comes to believe the truth, but the interference of another agent, such as a "helpful demon." Here, as in other cases, credit for true belief and hence knowledge is undermined. And, here again, this seems to be a general phenomenon regarding credit. So, for example, the influence of a helpful demon would undermine moral credit as well. I thank Daniel Nolan and Tamar Gendler for raising this kind of concern.

^{xxxiv} Zagzebski raises the problem in *Virtues of the Mind*, pp. 300-2, and in a more extended way in "From Reliabilism to Virtue Epistemology," *Proceedings of the Twentieth World Congress of Philosophy, Volume V: Epistemology* (1999): 173-179. Reprinted in expanded form in Axtell, ed., *Knowledge, Belief and Character: Readings in Virtue Epistemology* (New York: Rowman and Littlefield Publishers, 2000). See also Zagzebski's contribution to this volume.

^{xxxv} Riggs takes a similar approach to the value problem. He writes, "When a true belief is achieved non-accidentally, the person derives epistemic credit for this that she would not be due had she only accidentally happened upon a true belief. . . . The difference that makes a *value* difference here is the variation in the degree to which a person's abilities, powers, and skills are causally responsible for the outcome, believing truly that *p*." See "Reliability and the Value of Knowledge."

^{xxxvi} I would like to thank Robert Audi, Heather Battaly, Michael Bergmann, Stewart Cohen, Keith DeRose, Tamar Gendler, Stephen Grimm, Daniel Nolan, Philip Quinn, Wayne Riggs, Ted Sider, Eleonore Stump, Ernest Sosa, Fritz Warfield, and Linda Zagzebski for their helpful comments in discussion and on earlier versions of the paper.