

Risk Sensitive Credit

Maura Priest¹

Received: 15 December 2016 / Accepted: 1 February 2018 / Published online: 21 February 2018
© Springer Science+Business Media B.V., part of Springer Nature 2018

Abstract Credit theorists claim to explain the incompatibility of luck and knowledge and also what makes knowledge valuable. If the theory works as well as they think, it accomplishes a lot. Unsurprisingly, however, some epistemologists remain unsure. Jennifer Lackey, for instance, proposes a dilemma that suggests credit theories are either too strong or too weak. Her criticism has been hard to overcome. This paper suggests a modified account of knowledge as credit for true belief that allows credit theorists to better counter Lackey’s criticism. I call my version of credit theory Risk Sensitive Credit. Under my account, an agent deserves credit just in case she believes truly on account of her reasonably accurate epistemic risk assessment. This assessment need not include higher order beliefs or even enter into conscious thought. Recent work in cognitive science, for instance, suggests that our visual faculties, in the absence of our direct awareness, work in accordance with a risk sensitive framework. This research will be referenced to help explain the dynamics of barn facade Gettier cases.

1 Credit and Credit Theory

As much disagreement as there is among epistemologists, there are a few points on which most can come together. We think that knowledge is valuable, that it is incompatible with certain kinds of luck, and that finding out just what it takes to know something would be pretty nice. Credit theory is one of the latest schools of epistemic thought that attempts to address all these issues we care about. What it takes to know, credit theorists argue, is creditworthy true belief. There is

✉ Maura Priest
mp3588@columbia.edu

¹ Radford University, 901 McGee Street Apt. 402, Kansas City, MO 64106, USA

disagreement over qualifications for creditworthiness, but at minimum, the true belief must be attributable to the believer.¹ This much alone gives epistemologists a lot. It explains what kind of luck is incompatible with knowledge and why lucky belief is a problem at all. Knowledge sabotaging luck interferes with, prevents, or nullifies an agent's creditworthiness; instead of attributing a true belief to the believer, we are inclined to say, "He just got lucky."

Creditworthiness is normative: if S deserves credit, S did what she epistemically ought to have done. John Greco, for instance, argues that, "a central task of epistemology is to provide an account of the normativity involved" (2010: 4). Whatever the specifics, normativity seems linked to ability. To return to Greco, knowledge is, "a kind of success through ability" (2010: 104). Ernest Sosa puts things in terms of competence: "A performance is apt if its success manifests a competence seated in the agent" (2009: 12). Although their views are not identical, both Greco and Sosa claim that knowledge acquisition demands (1) success (true belief) and (2) that the agent herself contributed to (1). The second component not only ensures a belief is owing to the agent instead of luck, it explains why knowledge is valuable. Because epistemic credit is an achievement and we value achievements, we also value knowledge.

So credit theorists claim to explain why luck and knowledge are incompatible and also what makes knowledge valuable. If the theory works as well as they think, it accomplishes a lot. Unsurprisingly, however, some epistemologists remain unsure. The most compelling objection is perhaps a dilemma posed by Jennifer Lackey. She argues that credit theories are either too strong or too weak; that the framework will either grant knowledge when it should deny it or deny knowledge when it should grant it. This paper suggests a modification to Greco's account that allows credit theorists to better counter Lackey's criticism. With this change, credit theory can rival any conceptual account on the epistemic market.

I call my version of credit theory Risk Sensitive Credit (RSC). Here is a quick and dirty definition: an agent deserves credit just in case she believes truly on account of her reasonably accurate epistemic risk assessment. This assessment need not include higher order beliefs or even enter into conscious thought. Recent work in cognitive science, for instance, suggests that our visual faculties, in the absence of our direct awareness, work in accordance with a risk sensitive framework. This research will be referenced to help explain the dynamics of barn facade examples. In preview, Henry's mistake can be understood as an instance in which his perceptual system's risk framework goes awry.

In Sect. 2, I disambiguate two senses of credit and argue that the failure to distinguish them has created unnecessary confusion. Section 3 discusses Lackey's challenge regarding barn facade examples. Section 4 reviews Greco's response to Lackey's dilemma and argues that it is inadequate. Section 5 argues for a modification to Greco's account, drawing on ideas from Sosa, David Henderson and

¹ Another term that has been used is "creditable." However, as Greco has noted, "'[C]reditable' is ambiguous between 'praiseworthy' and 'attributable'" (2012: 5). Credit for true belief need not imply praiseworthiness for true belief.

Terry Horgan. With reference to the previous section, Sect. 6 resolves Lackey's dilemma.

2 Confusion in Chicago

Lackey's challenge case concerns an agent who easily acquires testimonial knowledge. Let us revisit the example from her original article:

CHICAGO VISITOR: Having just arrived at the train station in Chicago, Morris wishes to obtain directions to the Sears Tower. He looks around, approaches the first adult passer-by that he sees, and asks how to get to his desired destination. The passer-by, who happens to be a Chicago resident who knows the city extraordinarily well, provides Morris with impeccable directions to the Sears Tower by telling him that it is located two blocks east of the train station. Morris unhesitatingly forms the corresponding true belief (2007: 352).

The criticism that the above example aims to draw out is this: it seems that because of the stranger's testimony, Morris knows the way to the Sears Tower. Nonetheless, it appears that Morris is unworthy of credit. All he did was ask a random stranger, which takes little cognitive work. If anything, the testifier deserves the credit.

The Morris dispute brings out two confusions. The first involves "shared credit," and the second the relation between credit and effort. The literature has adequately addressed the former: we have no reason to think that the testifier's creditworthiness in any way suggests the recipient's unworthiness. Credit can be shared in such a way that each party deserves "significant" credit for the given instance of knowledge acquisition. Hence critics who point out the testifier's creditworthiness speak to something which is true but irrelevant. It is true that the testifier deserves a lot of credit for Morris's true belief. But this is irrelevant to the question of Morris's creditworthiness. Because credit can be shared, the testifier might deserve a lot of credit while Morris is still creditworthy to the extent needed for him to acquire knowledge. As Greco has argued, "[C]redit for an achievement, gained in cooperation with others, is not swamped by the able performance of others. It's not even swamped by the outstanding performance of others. So long as one's own efforts and abilities are appropriately involved, one deserves credit for the achievement" (2009: 228). In other words, what matters is not *quantity* but *quality*.² Creditworthiness should be measured on a scale of epistemic excellence, not relative epistemic contribution, and not, as I shall argue, epistemic effort.

As mentioned, some credit theorists have argued that credit can be "easy." Said differently, some credit theorists have argued that the credit needed for knowledge need not be credit for doing something epistemically amazing or something of great epistemic difficulty. Rather, credit can be earned for doing, in the words of Wayne

² "The new picture retains the idea that S's abilities are important in the causal story regarding how S comes to have a true belief, but it makes 'importance' qualitative rather than quantitative. Rather than saying that S's abilities must be 'important enough' in the causal story, we will say that they must be 'important in the right way'" (Greco 2012: 13).

Riggs, “simply...enough.” I agree with this sentiment. Sometimes what is enough for knowledge is neither epistemically difficult nor epistemically impressive. However, I also think this idea can use further delineation. Critics of credit theory may be skeptical of this “easy credit,” because these critics are using a different definition of “creditworthiness” than are credit theory advocates. Let me explain.

CHIACGO VISITOR brings to light divided pre-theoretical intuitions concerning creditworthiness. On the one hand, we see creditworthiness tied to admirable or even ethical characteristics. We think that the hard-working student with little ability deserves more credit than the gifted lazy student. Notwithstanding, we might give the former student B’s and the latter A’s. So in this other sense the lazy student deserves more credit. This is demonstrated in the higher grade: he earned more academic credit, i.e., an “A” is worth more academic credit points than a “B.” We conceptualize credit in terms of both effort and excellence, which often go hand in hand, but often come apart. Indeed, at times our standards of excellence are completely divorced from effort.

When criticizing credit theory, Lackey seems to assume the effort conception of credit. In her own words, “[T]he absolutely *minimal work* being done by the recipient of testimony casts serious doubt on the plausibility of him deserving credit for the truth of his belief” (2009: 37) (emphasis added).³ In another paper Lackey compares the testimonial case of Morris to a vision case. Her words are as follows, “For even in the simplest and most *effortless* cases of perceptual knowledge...” Given that Lackey is comparing this case to the Morris example, here Lackey implies that Morris’s acquisition of knowledge via testimony was also “effortless.” However, we should note that references to a “minimal workload” or “effortless” knowledge acquisition only cast doubt on Morris’s creditworthiness IF we understand credit in terms of effort. If, however, we are conceptualizing credit in terms of excellence, then “minimal work” or epistemic tasks that are “effortless” counts neither for nor against creditworthiness.⁴ And this is how credit theorists seem and ought to understand the concept. Credit theories are virtue theories and virtue theories are theories of excellence. As Greco has noted, “...the notion of ‘virtue’ in play is *personal level excellence*” (2012: 9) (original emphasis). Yes, knowledge is a success term. But often to our dismay, success and effort come apart. As a former track coach, I have seen this too many times. Consider what will be the

³ At another point Lackey again puts her dilemma in terms of the “minimal work” being done by the testimony recipient, “As should be clear, either horn of this dilemma undermines the Credit View of Knowledge at its core. For, on the first horn, credit may be adequately blocked in Gettier-type cases, but only at the expense of also blocking credit in countless cases where testimonial knowledge is intuitively present despite *minimal work* being done on the part of the hearer. And, on the second horn, credit is secured in cases of testimonial knowledge where such *minimal work* is done by the recipient, but only at the expense of also securing credit in Gettier-type cases. Either way, the Credit View not only fails to shed light on what is absent in Gettier-type cases, but it also fails to explain the additional value that knowledge has over merely accidentally true belief” (2009: 34) (my emphasis).

⁴ I believe there is enough textual evidence to claim that Lackey was indeed thinking of credit in terms of excellence. However, I admit that although work and effort are often connected they are not always. So it is possible that Lackey’s references to minimal work were not meant to capture an effort conception of credit. In the end, it seems nonetheless important to clarify what kind of credit we are working with for the purpose of resolving Lackey’s dilemma.

first of many sport analogies: A talented sprinter might blow by his competition with less effort than it takes most people to tie their shoes. For this natural athlete, success was achieved with hardly any effort at all. Another athlete might train hours a day and have no shot at out racing the effortless excellence achieved by the natural athlete.

Because the credit Morris earned and needed was credit in terms of excellence, whether he worked hard or had significant help is irrelevant. Let us again return to Riggs and quote him more completely, “There is not some stable threshold of effort...that must be superseded... (We) simply have to do enough to bring it about.” (2009: 218). Indeed. And sometimes very little is also just enough. CHICAGO VISITOR suggests that epistemic excellence and epistemic laziness are not mutually exclusive. Consider a variant example:

BOOTSTRAP BUCK: Buck is a rugged epistemic individual, a “pull yourself up from your bootstraps” type of guy. He steps off the Chicago train determined to form a true belief about the Sears Tower: on his own. He wants full epistemic credit. Accordingly, he refuses to ask for directions, turns off his smart phone and will not even glance at a map. Instead, Buck uses intuition, the location of the sun, memory, and wholesome epistemic grit. Thirteen hours later, exhausted but satisfied, Buck finds himself in front of the Sears Tower. With tired pride, he forms the justified true belief, “The Sears Tower is located at ‘233 S. Wacker Drive.’”

We see that no one other than Buck deserves credit for his true belief. But is he Morris’s epistemic superior? Probably not. Maybe Buck had non-epistemic reasons to behave as he did. Perhaps he wanted the exercise, was seeking adventure, or preparing for a survivalist TV show. But if Buck’s goals were purely epistemic, it seems he wasted a lot of time. The effort spent searching could have been used on other epistemic endeavors: reading Descartes Meditations, conversing with a physicist, or listening to a lecture at Northwestern University. Whatever the alternative, it would surely offer greater epistemic returns than the tower expedition.

Using more cognitive prowess than necessary may contribute to success in an instance, but insofar as extraneous effort impedes an epistemic life well lived, it is not an epistemic virtue. At the very least, there is no reason to think Buck’s behavior was preferable or more virtuous than Morris. But if Buck did not exemplify a virtue, must credit theorists deny that he acquires knowledge? This would be a serious problem, for regardless of the inefficiency, it does seem Buck knows when he finally arrives at his destination. Credit theorists can respond as follows. Credit should be understood in two senses: diachronic credit and synchronic credit. The former is acquired when an agent’s belief forming mechanism (1) is attributable to the agent, and (2) contributes to an epistemic life well-lived over time. The latter is reserved for belief forming mechanisms which (1) are attributable to the agent, and (2) reliably lead to a true belief in the particular situation at hand. Each kind of credit can be had without the other. An agent might acquire a belief in a given instance that makes him worthy of synchronic credit even though his form of acquisition was not wise from the standpoint of an epistemic life well-lived. Arguably, this is what happens with Buck. So Buck is not worthy of diachronic credit but he does earn

synchronic credit and the latter is what counts for knowledge. On the other hand, an agent might do everything right and yet believe falsely. I will argue that this is what happens with Henry and even though he might be worthy of a type of diachronic credit, the type of credit which turns a true belief into knowledge is synchronic. Henry fails to earn synchronic credit. For the rest of the paper I will focus exclusively on synchronic credit.

With the distinction between these two types of credit in mind, let us revisit the question of Morris's creditworthiness. In judging whether an agent is worthy of synchronic credit, we must consider at least three things. First, was the means of belief acquisition reliable? Second, can we attribute the use of such reliable means to the believer? Third, can we identify the presence of any credit undermining defeaters? Creditworthy beliefs will have a "yes" answer to the first two questions and a "no" to the third. Regarding question one, many do have the intuition that asking a competent looking stranger for directions reliably (or reliably enough) leads to true belief. We seem to think at least, that if the passerby is incompetent she will admit so, rather than just make something up at random.⁵ After all, if directional inquires usually failed, there would be no point in asking. So part of Morris's epistemic excellence is just this excellence of reliability. Think of a soccer player that kicks the ball into a hard to reach corner of the net. This is an act of athletic excellence insofar as this means of shooting reliably leads to scored goals. Likewise, then, asking a stranger for directions is an act of epistemic excellence insofar as this means of belief formation reliably leads to true beliefs.

Next we must consider Morris's attribution. Assuming Morris is an ordinary fellow, he has some past experience with directional inquiries, or has observed others, or has conversed about the practice.⁶ Hence what is attributable to Morris is the non-accidental undertaking of requesting directions alongside a personal history of either success with this practice or acquaintance with the success of others. This attributable is a source of excellence and can explain Morris's creditworthiness.

The last step in assessing Morris's creditworthiness is a typical search for defeaters. We are stipulating that the person Morris asked was not a drunk, a child, or a squirrel. Let us also presume no one nudged Morris on the train, whispering,

⁵ In "Knowledge and Credit," Lackey seems doubtful of this reasoning. In her own words, "[W]hat does an honest, competent passerby look like that would enable Morris to reliably distinguish her from a dishonest or directionally- challenged one... on any plausible reading of CHICAGO VISITOR, Morris could have just as easily approached a competent-looking liar or a directionally challenged speaker as he did an honest, knowledgeable, Chicago resident when asking for assistance in finding the Sears Tower" (2009: 31–32). If by "'just as easily,'" Lackey means to imply "just as likely," I think this at least needs more argument. My experience with asking for directions suggests otherwise. And it seems plausible that our intuition that we can gain knowledge by asking for directions is linked to our experience of this being a fairly reliable way of finding a desired location. On the other hand, if by "just as easily," Lackey means something other than likelihood, I will argue (in Sect. 5) that likelihood is what matters.

⁶ Some might argue that Morris could acquire knowledge even if he was unaware of the practice of requesting directions and did ask on a whim. A case like this would be harder, although perhaps not impossible, for credit theorists to explain. I do not think, however, that credit theorists would need to account for this unusual interpretation. The force of Lackey's point is grounded in the ordinary intuition that Morris can attain knowledge by asking a stranger. The ordinary background story of such an occurrence is one in which Morris is familiar with the practice of asking for directions. The intuitive pull of her counterexample then, is linked to an ordinary construal of the case.

“Whatever you do, don’t ask for directions in Chicago!” It is also crucial that Morris did not have a lie-detecting smart-phone app that was set off when the passerby begun to speak. If all this is the case and there are no other disqualifying defeaters, then we have no reason to strip Morris of credit. Morris, and many others who gain knowledge via testimony, act with epistemic excellence insofar as they use a reliable means of true belief acquisition and such use is attributable to them. In Sect. 5, we will return to testimony, and I will go into further detail about the excellences of testimonial recipients.

3 Lackey’s Dilemma: Part 2

Suppose we understand credit in terms of excellence, and that Morris thereby acquires credit and knowledge. This responds to the first horn of Lackey’s dilemma; it explains why Morris is creditworthy. What about the second horn; does understanding credit in terms of excellence make the theory too weak? Let us revisit the example Lackey uses to support her weakness accusation.

FAKE BARN: Henry is driving through the country, looks out the window, and forms the belief, “That’s a barn.” His belief is true. However, Henry is in Fake Barn Country, and the barn he saw was surrounded by fakes.⁷

We do not want to attribute knowledge to Henry. But Lackey argues that credit theories must do just this to avoid inconsistencies. Compare Henry and Morris: both have true, justified, beliefs attributable to their own agency. Lackey concludes that either both Morris and Henry deserve credit or neither do. If both deserve credit, then against intuitions, Henry knows that he is looking at a barn. If neither, then against intuitions, Morris fails to know the Tower’s location. Explaining Morris’s creditworthiness solves one problem at the expense of another: credit theorists must now explain why Henry *lacks credit*.

Prima facie, conceptualizing credit in terms of excellence is little help with dilemma part II. If Morris can be credited with epistemic excellence, it seems so can Henry. Morris earns credit because he utilized his abilities in the right sort of way, a way that is attributable to him and that reliably leads to true belief. And credit theory critics will argue that we can say the same of Henry. After all, Henry’s true belief is at least partly attributable to his reliable vision. In Sect. 5, I argue that contrary to first appearance, credit theorists can show that Morris is creditworthy but Henry is not. First, however, I want to explain why current credit theories, Greco’s theory in particular, falls short of offering a fully satisfying account.

⁷ This example is based off of one credited to Carl Ginet and cited by Goldman (1976).

4 Greco's Solution

4.1 The Argument

According to Greco, “an agent might have an ability relative to one environment but not another” (2012: 42). In this fashion Greco argues that we might distinguish CHAICAGO VISATOR and FAKE BARN. In Greco's words once again, “Henry believes from a disposition that is reliable relative to normal environments, but not relative to the environment he is in. Accordingly, Henry does not know that the object he sees is a barn” (2012: 25) Formally, Greco argues as follows:

1. Henry's perceptual disposition regarding barns is not reliable relative to Fake Barn Country (and therefore does not count as an ability).
2. Credit and knowledge are only acquired if belief is produced by a disposition that is reliable relative to the environment (that is, produced by an ability).
3. Therefore Henry acquires neither credit nor knowledge.

In contrast to Henry, Greco argues that Morris's relevant disposition *was* reliable on that busy Chicago sidewalk and so rendered Morris creditworthy. Greco's strategy of linking ability to environment may seem innocent enough, but its true plausibility depends on how we circumscribe environments. To this end, Greco offers an analogy to illustrate which environments *exclude* creditworthiness: “(Derek) Jeter has the ability to hit baseballs in typical baseball environments, but presumably not in an active war zone, where he would be too distracted” (2012: 42). Just as Jeter's baseball abilities are relative to a peaceful environment, Henry's abilities are relative to a traditional farm environment. And this, Greco suggests, solves the second horn of Lackey's dilemma. Because the environment in Fake Barn Country is unusual, and in a way that makes Henry unreliable, he lacks the relevant ability and so acquires neither credit nor knowledge. Morris's environment allows him to form his belief reliably, and so he can come to know the Tower's location.

4.2 Risk Versus Reliability

A potential difficulty with Greco's response is the apparent conflation of risk and reliability. An example from Duncan Pritchard (2012) helps illustrate. I paraphrase:

PIANO: A pianist is performing in a threatening environment: he is surrounded by walls which could collapse and flood the room with water. But as long as the walls remain intact, he performs excellently.

Pritchard's point is that despite of the risky environment, the pianist properly exercises his abilities. Similarly, there is a sense in which Henry's visual ability works just fine in Fake Barn Country. What I am referencing is the physiological functioning of his human eyesight. This is a reference not to epistemic reliabilism, but simply physiological functioning of a sighted human being. Consider, for instance, if Henry were to go to an optometrist and get his eyes tested. Or suppose an optometrist was right there in Fake Barn Country. If Henry was tested, he would

pass with flying colors (or so the structure of the example leads us to assume). So although Henry is in a risky situation, he has good vision and a clear view of a large dry-good from a modest distance. In this sense, his abilities are performing just as they ought. We can modify one of Greco's examples and make a similar point.

HOMERUN: Jeter is playing baseball in a war zone for a "support the troops" charity game. Unexpectedly, chaos erupts. However, Jeter's love of the game compels him to continue. He receives the perfect pitch and the ball flies past enemy fire and clear over the makeshift stadium wall.

It seems that despite the risky environment, the hit is still attributable to Jeter. In a regular game Jeter is creditworthy insofar as his coordination, strength, and determination contribute to his successful hit. And these are the same features that lead to success in HOMERUN. If he had been too distracted and struck out, this need not count against his abilities as it would in Yankee Stadium. As Greco says, "[I]t does not count against Derek Jeter's ability to hit baseballs that he would fail in poor lighting conditions." (2012: 42). Fair enough. But when Jeter *does* succeed in poor lighting conditions, or in a war zone, we might still attribute the success to ability. There is a sense in which HOMERUN Jeter gets lucky. But it is neither luck that renders his success accidental nor luck that disqualifies him from athletic credit. Disqualifying luck contributes to success and thereby weakens or eliminates an agent's own contribution. But Jeter used his abilities in the same way as he would in a professional baseball stadium. He was lucky nothing interfered with his hit. But since he was lucky, he avoided the dangers and his abilities secured success. Consider, for example, if I was in the same situation as Jeter. No matter what happened in this environment, I would not hit a home run because I lack the underlying baseball abilities in the first place. It is only because Jeter has these abilities to begin with that he is able to succeed in a war zone and hit a homerun.

I believe Pritchard and I are trying to make the same point: environments can be the wrong type (relative to ability) in two distinct ways:

- (1) The environment can be such where the ability in question does not work at all, or is simply inapplicable to the environment.
- (2) The environment can be one in which the risk that something interferes with the ability is very high.⁸

An example of the first type might be playing a piano underwater. When we talk about someone having piano playing skills, we do not consider underwater playing part of that ability. In fact, we assume that "playing the piano" strictly does not include playing underwater. An example of the second type is an environment that is at high risk of suddenly becoming underwater. In this case, playing the piano is a relevant ability and an agent can succeed in playing the piano: there is just a high

⁸ Butts (2014) discusses many issues related to Greco's construal of abilities. Butts argues that Greco is better off defining abilities in the same way as Sosa, Kallestrup, and Pritchard. I am arguing that there is an inconsistency with the way Greco uses his own conception of abilities. The inconsistency is argued for clearly in Sect. 4.3.

risk that things go wrong. When they do, the ability is no longer an ability. (i.e., once you are underwater, we do not expect you to be able to play the piano). If things do not go wrong, however, you get lucky and your ability to play piano is still exercised.

The piano comparison is relevant because Greco seems to suggest the environmental problem with Henry is (1) when it is really (2). Henry does not have the perceptual ability to correctly identify a fake barn. But this is not relevant as long as he is looking at a real one. And what happens in Fake Barn Country is Henry gets lucky: he is at a high risk of his perceptual ability being interfered with but this does not end up happening. Because Henry gets lucky, we cannot simply say he lacked the ability in the environment. He did not lack the ability, there was just a high risk that the conditions would change (i.e., his eyes would fall upon a fake) and if they did then he would lack the ability to identify the object via sight. Lucky for Henry, conditions did not change. The point is that Henry indeed exemplified visual abilities in the fake barn environment.

4.3 Informational Needs

Because Greco argues that solving barn cases involves understanding ability relative to environment, not only must we identify the relative environment, but we also must identify the ability. To put things simply, Greco could offer a general or a specific understanding of abilities. But both options have problems. The problem with a general construal of abilities is that this leads to the wrong results in Fake Barn country. For if we consider Henry's relevant ability "eyesight" or "vision," then it is easy to make the case that Henry *does* properly utilize the relevant ability. After all, the eyesight Henry uses is the same eyesight that would pass an eye exam at the ophthalmologist's office. This result is not what we want, for it is counterintuitive that Henry has knowledge (and Greco himself agrees it is counterintuitive). What Greco needs is a construal of abilities that is more specific and hence disqualifies Henry from properly exercising them in the fake barn environment (which would thus, exclude Henry from acquiring knowledge). Specific constructions might include, "The ability to identify fake barns" or "The ability to distinguish fake barns from real ones." This specific construal of abilities does allow us to say that Henry lacks the abilities in the environment and hence does not acquire knowledge. The problem with this construal, however, is that it is inconsistent with how Greco defines abilities in other examples. Consider the following example involving a Grizzly bear, which I paraphrase below⁹:

GRIZZLY*: Timothy walks into a cave and is face to face with a hungry Grizzly. He forms the true belief, 'I am face to face with a Grizzly bear'. Immediately after aforementioned belief formation, Timothy is eaten by said Grizzly.

The original justification for bringing up the Grizzly bear example is not relevant for our purposes. What is relevant is how Greco defines Timothy's epistemic abilities.

⁹ Greco (2012) credits the original example to Joe Salerno.

Here is Greco in his own words, “Presumably, the ‘sort’ of ability in question is visual perception, and the ‘way’ in question involves the normal exercise of that ability in normal enough lighting conditions” (2012: 23–24). So Greco suggests construing abilities generally. We should think of Timothy’s ability as “perception in normal enough lighting conditions,” not “perception in the presence of a hungry bear.” However, let us contrast GRIZZLY* with FAKE BARN. There are various ways we might define Henry’s abilities. The option, however, that is analogous to the bear example is, “perception in normal enough lighting conditions.” Yet defining abilities in this way means that Henry indeed exercised a reliable ability and hence should acquire knowledge. This is exactly the result that Greco wants to avoid.

To be clear, the point I am making is not a generality problem of environments but one of abilities. My criticism would not be addressed even if there was no controversy whatsoever regarding environmental scope. For once an environment is identified, we must then determine the *ability* in respect to that environment. Suppose we agree that Henry’s environment is Fake Barn Country. We must then determine whether the ability in question is simply “vision,” or “skill in distinguishing real barns from fakes” or “perception in a barn environment.” The most obvious way to circumscribe Henry’s ability is simply “vision” or “perception.”¹⁰ If abilities are relative to environment, then “vision (or perception) in Fake Barn Country.” But Henry’s perception in Fake Barn Country *is* reliable. After all, he can acquire perceptual knowledge about clouds, street signs, lamas, etc. It seems forced to claim his perception is unreliable only regarding barns. Lackey has already recognized as much, and in an important footnote argued the following:

While Greco may be right that reliability is relative to an environment, it is unclear why he thinks that Henry’s perception is not reliable in the example under consideration. For surely Henry would form mostly true beliefs by relying on perception in the environment in question, e.g., he would form true beliefs about farmers, horses, pigs, trees, grass and so on. The only sense in which his perception is not reliable in the relevant environment is with respect to distinguishing real barns from barn façades while driving in his car past them. But individuating cognitive faculties this narrowly leaves the door wide open to worries about the generality problem (2007: 355).

How can Greco answer the above concern? In response to generality accusations, he argues that scope must be understood in accordance with practical needs.¹¹ It is unclear, however, how this solves the problem. Suppose Henry had no practical need to distinguish fake barns from real ones. Let us stipulate he was driving right through the country on the way to New York; identifying barns would further no

¹⁰ Perception can be understood to encompass a much wider range of faculties than just vision. However, for my purposes here, the two terms will be used interchangeably.

¹¹ Consider the following from Greco, “[W]e can solve this (generality) problem by making use of an idea that we invoked above: that the concept of knowledge serves the purposes of practical reasoning. Specifically, that the concept of knowledge is used to flag good information and good sources of information for use in practical reasoning. If this is right, then we have an answer to the generality problem: relevant parameters should be specified according to the interests and purposes of relevant practical reasoning” (2010: 78).

practical purpose. Notice that even in this case, we would not want to say that Henry knows.

4.4 Using Abilities Versus *Attribution to Abilities*

So at this point I have made a strong case for why Henry has properly functional visual abilities in Fake Barn Country. It is also clear that these properly functioning visual abilities play a role in Henry arriving at his true barn belief. However, one might still attempt to defend Greco by arguing that although these abilities are indeed functioning, and although they indeed play a role in Henry arriving at his true belief, it does not follow that Henry's true belief is *attributable* to these abilities.

What exactly does it mean for a true belief to be attributed to one's abilities? Well, Greco has this to say about the matter:

The present account employs a couple of plausible assumptions. The first is that...an outcome is produced by means of multiple contributing causal factors. The second is that explanatory salience distributes unevenly... For example, suppose that sparks cause a fire...the sparks do not cause the fire all by themselves—there has to be oxygen present, as well as combustible material, etc. Nevertheless, in the typical case it will be true that the fire is attributable to the sparks. That is because, in the typical case, the sparks *explain why* the fire started. On the other hand...the presence of oxygen in a warehouse does not explain why a fire started...(2012: 44).

So it seems that Greco is arguing for a type of “explanatory salience” in a complex explanatory chain. (or at least, “salient enough”). Hence a true belief is, “due to” or “attributed to” X, if and only if X is a particularly salient factor in this explanatory chain. Regarding abilities then, the abilities must be manifest such that it makes sense to say that S's abilities are a particularly salient factor in arriving at the relevant true belief. In other words, it makes sense to say that S's true belief is because of S's abilities.

Now Lackey's challenge to Greco, is that he somehow parse a theory that can account for why Morris's true belief is attributable to abilities, but yet Henry's true belief is not. Said differently, Greco must propose a theory with the following results: Morris's abilities ARE a salient part of the explanation of his true belief. However, Henry's abilities ARE NOT a salient part of the explanation of his true belief. The problem, as I see it, is that this (without further details) remains a fine line to draw between the Henry example where his true belief is NOT because of his abilities, and the Morris example where the true belief is because of abilities. My point is that as things are with Greco's account, there are two ways to read the Fake Barn Country example. One of my claims is that the line between these two ways of thinking is not a thick line. Indeed it is especially thin. Here are the two possibilities:

- (1) Henry forms a true belief due to his properly functioning visual abilities.
- (2) Henry forms a true belief by using his properly functioning visual abilities, but NOT “due to” these abilities.

To make my paper worthwhile, I do not believe I need to show that (2) is false and (1) is true. I simply think it must be the case that (1) is plausible, or that there is a fine line between (1) and (2). Because if (1) is at all plausible, there is value in arriving at an alternative theory that can approach the challenges posed to credit theory in a different and yet compelling fashion.

What I am trying to do is offer a modified credit-theory that is an alternative to Greco's (and Sosa's, and others) because there are spots where this theory has shaky responses to serious objections. One such spot is the fine line between (1) and (2) above. What I am about to propose is what I consider a promising alternative way to slice it, i.e., a way that might have extra appeal to those who find Greco and Sosa's answer to Lackey's dilemma unsatisfying. In spite of initial appearances, I will argue that by delineating creditworthiness in greater detail, credit theorists can adequately respond to each of the two challenges posed by Lackey's dilemma.

5 Risk Sensitive Credit

5.1 Risk Sensitivity

For any proposition p , an agent might believe p , believe not- p , or withhold belief. Believing can be epistemically risky. What I mean is that few beliefs we hold with epistemic certainty. Because we are rarely certain of our beliefs, going forward from a space of withholding belief to believing is a risk. (At least, it is a risk for those of us who want to avoid believing falsehoods.) At times, the risk of false belief might not be worth the potential reward. Keeping this in mind, S 's belief is what I call *risk sensitive* only if the likelihood of false belief is low enough that belief is the best epistemic option. Said differently, the risk ratio must be such that the expected epistemic value of believing p is greater than the expected epistemic value of withholding belief (assuming that a true belief has positive epistemic value and false belief negative epistemic value).

Sosa has discussed ideas that align well with my above thoughts about risk sensitivity. In his own words, “[One’s] meta-competence governs whether or not one should form a belief at all on the question at issue, or should rather withhold belief altogether” (2009: 14). Elsewhere Sosa argues that “A performance can thus easily fail to be ‘meta- apt’, because the agent handles risk poorly, either by taking too much or by taking too little. The agent may fail to perceive the risk, when he should be more perceptive; or he may respond to the perceived risk with either foolhardiness or cowardice...” (2009: 12).

What I call risk sensitivity is similar to Sosa's meta-aptness. S 's belief lacks risk sensitivity if she takes “too much risk or too little.” Going overboard in epistemic riskiness results in a belief in p despite of the high chance of p 's falsity.¹² On my account, an agent only earns credit if her belief is risk sensitive, therefore, taking too

¹² Alternatively, she might believe p when there is a low risk of falsity, but she will still lack credit if her own abilities do not assess this risk. In other words, she will lack credit if her epistemic success is not due to her risk assessment abilities.

much risk precludes one from earning epistemic credit. Because credit is required for knowledge, taking too much epistemic risk precludes knowledge acquisition.

In addition to the epistemic mistake of taking too little risk, an agent might handle risk in the opposite fashion. At times, an agent might take too little epistemic risk. Taking too little risk is withholding belief in p despite of the low chance of p 's falsity. The bad epistemic consequence in taking too little risk is different than the consequences from taking too much risk. In the latter instance, an agent acquires a false or unjustified belief. However, when an agent takes too little risk she incurs an epistemic opportunity cost: she fails to have a true belief that she would have if she had only been a little more open to risk. She was in a situation in which her information merited belief but she still withheld. Unnecessarily withholding belief is a failure of epistemic excellence and this failure explains why someone who takes too little risk earns no epistemic credit.

When an agent does not err in either taking too much risk or too little, she sometimes acquires a risk sensitive belief. The first step in acquiring a risk sensitive belief is assessing epistemic risk and in so doing avoiding the twin downfalls of risking too much or too little. Now let us be clear that while epistemic credit demands risk assessment, according to my account this risk assessment need not include any higher order beliefs or even the possibility thereof. (If "assessment" sounds too reflective, you may prefer to think of risk "accommodation"). Now it might appear that risk assessment without the possibility of higher order belief would not concur with Sosa's understanding of meta-aptness which does require higher order belief.¹³ At this point we can turn to David Henderson and Terry Horgan. Henderson and Horgan question Sosa's theory:

We ourselves find very plausible the idea that competent risk assessment, as an aspect of the process of forming a belief, is required in order for that belief to constitute fully human knowledge. But we doubt whether such competence needs to take the form of a higher-order belief; and we also doubt whether a first-order belief can qualify as any kind of knowledge if is formed in a way that *utterly lacks* the aspect of competent risk assessment (2013: 601) (original emphasis).

The risk sensitivity I advocate aligns with Henderson and Horgan on both counts: S cannot know p unless S (or S's abilities or S's cognitive system) assessed (or accommodated) p 's risk, but this can take place without higher order belief. Risk assessment itself need not involve any type of belief whatsoever. Notwithstanding, risk assessment is necessary for knowledge of any kind. (The assessment itself is what conveys information to the agent regarding the risk involved in believing p . Conveying this information, however, need not come in the form of a belief.)

Let us return to Sosa and his account of risk assessment. I may have been too quick in concluding that Sosa's account rules out the type of minimally reflective

¹³ In Sosa's own words, "We can now see that knowing something full well requires that one have animal and reflective knowledge of it, but also that one know it with full aptness. It requires, that is to say, that the correctness of one's first-order belief manifest not only the animal, first-order competences that reliably enough yield the correctness of the beliefs that they produce. One's first-order belief falls short if it is not appropriately *guided* by one's relevant meta-competence." (Sosa 2009: 16) (original emphasis).

risk assessment I advocate. In a reply to Henderson and Horgan's criticism, Sosa argues that what appears to be a disagreement between himself and Henderson and Horgan may only be a misunderstanding. Because Sosa understands "belief" generally, his notion of belief might include cognitive states which Henderson and Horgan were not including. This liberal understanding of belief leads to less stringent demands for risk assessment (given that Sosa's risk assessment demands higher order "belief"). The clarification brings Sosa closer to Henderson and Horgan's view, as well as closer to my own. All of us seem to agree that risk assessment can occur in the absence of highly reflective higher-order cognitive processing. Yet a few points of dispute might remain. Later in his reply to Henderson and Horgan, Sosa argues that at least some instances of knowledge might arise without risk assessment. He says, "...there can be pure animal knowledge with no admixture of risk- assessing reflection, no matter how implicit. A very basic sort of knowledge can be found *even below the level of animal knowledge...*since it is constituted by mere guesses rather than beliefs" (2013: 631) (original emphasis).

To motivate his point regarding this knowledge which is "below the level of animal knowledge," Sosa describes the experience of taking an eye exam. During an eye exam, there comes a point at which our vision is not sharp enough for confident assertion. We can, however, still see enough to guess. If sufficiently reliable, Sosa thinks that such guesses might constitute very basic knowledge, even though such knowledge includes neither belief nor risk assessment. Here is what I would say in response: if such guesses can qualify as knowledge, these guesses *must* involve risk assessment. Even at the point in the eye exam when we begin to call our answers "guesses," they are not what we would call *total* guesses. In other words, they are not the sort of guesses we would give if blindfolded. Rather, we make the kind of guesses that Sosa refers to on the basis of our less than perspicuous visual stimuli. "Guesses" based on such stimuli might involve risk assessment, for we can use this stimuli for just such purposes.

Let us turn to Sosa's comments concerning a 'super-blindsighter'. He says, "... a super-blindsighter can just find a belief within, despite having been guided to form it by no risk assessment...Can't this be knowledge that is purely animal and *entirely* unreflective, based on no risk assessment whatsoever" (2013: 629) (original emphasis). We see Sosa emphasize that super-blindsighter knowledge can be "entirely unreflective." The way I am using the term, however, a complete lack of reflection does not eliminate the possibility of risk assessment. It is unclear to me, then, why a super-blindsighter could not engage in below the surface risk assessment. Similarly with the regular blindsighter. For instance, Sosa claims that "[B]lindsighter beliefs derive...from subpersonal processes systematically reliable enough to yield a kind of knowledge" (2013: 629). As will be explained in Sect. 5.2, I agree that anything we might call blindsighter knowledge would derive from systematically reliable subpersonal processes. I would argue, however, that such subpersonal processes are reliable in virtue of their ability to assess epistemic risk. I suspect that Sosa and I simply have differing understandings of "risk assessment." According to my account, *all* instances of knowledge involve "risk assessment," where "assessment" is understood loosely enough that it might occur without any beliefs, implicit or explicit, and without any conscious awareness or reflective

processing whatsoever. The next section discusses visual studies that should further shed light on this possibility.

5.2 Cognition, Perception, and Risk Assessment

Henderson and Horgan suggest that, “[We] might have a trained capacity that manages to accommodate [risk] without articulation, automatically and quickly...” (2013: 603). While I agree with the aforementioned, we might also have innate cognitive capacities that *evolved* to accommodate risk. I suspect that Henderson and Horgan were thinking of “trained” loosely, and this was what they meant. In any case, visual studies confirm that automated cognitive processes can classify sensory data according to a risk sensitive framework. Consider the following commentary on a recent study,

...Bayesian concepts are transforming perception research by providing a rigorous mathematical framework for representing the physical and statistical properties of the environment... describing the tasks that perceptual systems are trying to perform, and deriving appropriate computational theories of how to perform those tasks, given the properties of the environment and the costs and benefits associate with different perceptual decisions (Geisler and Kersten 2002: 508).

The above suggests that perception works within a cost–benefit framework that balances the benefits of perceptual belief versus the risks. Further studies provide evidence that we update statistical frameworks according to perceived environment. In short, there is much more to perception than sensory data. To ensure accuracy, our perceptual system first receives sensory information, and then second and separately, accommodates this data in accordance with the environment and other circumstantial contingencies. Environmental awareness, combined with sensory input, leads to risk assessment. This again is supported with research in cognitive science:

[T]he objects that are likely to occur in a scene can be predicted probabilistically from natural scene categories that are encoded in human brain activity. This suggests that humans might use a probabilistic strategy to help infer the likely objects in a scene from fragmentary information available at any point in time (Stansbury et al. 2013: 1031).

Our perceptual system matches visual sensations to familiar objects given other information about the environment and contextual circumstance. Suppose you experience a visual stimulus of a small furry animal. If you believe you are in the forest, this stimuli might indicate a squirrel. Contrastingly, if you were at home, your unconscious cognitive processes might suggest that the animal is a cat. To earn epistemic credit and acquire perceptual knowledge, first, your sensory data must accurately (or accurately enough) reflect the perceptual object. In other words, your vision is not blurry, you are an appropriate distance from the object, and you are not under the influence of hallucinogens. If this holds, you have data to make a probability assessment in accordance with the environment and other relevant

conditions. To return to our earlier discussion, it seems possible that a blindsighter might engage in this type of calculation. Whether we are okay then attributing the blindsighter with knowledge will depend on various other factors, including how stringent we view the requirements for belief, or if we think belief is necessary at all. But that is a bit of a digression. Back to our visual studies:

[A]n ideal observer convolves the posterior distribution with a utility function (or loss function), which specifies the costs and benefits associated with the different possible errors in the perceptual decision. The result of this operation is the expected utility (or Bayes' risk) associated with each possible interpretation of the stimulus. Finally, the ideal observer picks the interpretation that has the maximum expected utility (Geisler and Kersten 2002: 508).

We can replace the ideal observer with the virtuous, or creditworthy, observer. Sensory input prompts the following evaluation: What are the chances that this stimulus comes from object O given environment E and circumstances C? The answer determines whether it is best to believe p, withhold belief, or believe not-p. Let us assume that a true belief is an epistemic benefit and a false belief a cost. Ideal agents believe p only if belief has the highest expected epistemic value. The creditworthy agent, which may fall short of the ideal one, believes p only if believing presents minimal epistemic risk. We can call this modification to existing credit theories Risk Sensitive Credit (RSC):

RSC: An agent's belief p is risk sensitive and hence creditworthy if (1) her own abilities assess belief risk, and (2) she correctly believes p because (1) indicates a reasonably low chance of p's falsity.

Some might object to the vagueness of "reasonably low." It is used for two reasons. First, it seems a fruitless effort to determine whether the risk of falsehood must be below 15, 10, or 5 percent. We might try to qualify things the way some more formal epistemologists do by saying that there must be less than a 2% risk of falsehood, for instance. Now not only is it easy to see why agreeing on the right amount is a near impossibility, but there might also be contextual circumstances. I would like to leave my theory at least open to contextualization, where the right level of risk might depend on the context. Now I say "leave open" because I also think it is possible to have a non-contextualized account of my view where the level of risk needed for justification must be constant across cases. When first introducing a theory, it is best to cast a wide net.

Second, philosophers who disagree about justificatory *degree* might still agree on justificatory *kind*. Hence philosophers might agree that an agent acquires knowledge if a belief is formed on the basis of a risk assessment that suggests a low chance of falsity. They might not agree on just how low that risk might be. We should also note that while it is hard to determine a threshold of epistemic risk, so is it hard to determine whether a process or agent is reliable or whether a close world is close enough. RSC is not unique in its vagueness. Rather than liabilities, we might see all these indeterminateness as theoretical virtues. At the borders, there is strong disagreement over whether beliefs qualify as knowledge. We might then expect that

any theory of knowledge which aligns with pre-theoretical intuitions will have borderline cases in which it is unclear whether believing p qualifies as knowing p .

I want to make one last point about my theory that may have been lost in the rest of the explanation. I have said that risk sensitive beliefs are managed by the proper assessment of informant relevant to the risk of falsity. So in order to have a risk sensitive belief, an agent first relies on her risk assessment capabilities to process information relevant to the truth or falsity of a given proposition. Such risk assessment need not itself be a metabelief. Said differently, simply because an agent has a creditworthy belief that p is true, it does not follow that she has a metabelief about the risk of p 's falsity. Our cognitive processes can assess belief risk sans metabelief. What matters is proper cognitive responsiveness to relevant information.

5.3 More on Risk

Risk sensitive belief is belief in accordance with reasonable risk assessment. What is risk assessment? Briefly, it is a means of analyzing and interpreting relevant data within an environment and set of conditions. Assessment goes about as follows: an agent's cognitive system, consciously or unconsciously, assesses the chances of p given what I call her "total information." Now total information *does not* refer to all the information there might be regarding a certain proposition, but simply all the information the agent happens to have. So an agent might have very little information about p , and hence this agent's "total information" would be sparse. In any case, such information consists of certain epistemic data D and epistemically relevant conditions C . That is, $P(D \& C)$. Risk assessment can go awry in at least three ways:

Risk Assessment Errors

- (1) Too much inaccurate data
- (2) Too much inaccuracy regarding the conditions
- (3) Seriously misinterpreting the meaning of the data given the conditions

Imagine a risk management company, SECURE is hired to assess the safety of a mansion hosting a prestigious fundraiser. SECURE might blunder through inaccurate data gathering, inaccurate conditional assessment, or misinterpretation of the data given the conditions. Examples of the first could include miscounting the fire alarms or misreading the thermostat. Either error would skew total assessment. But maybe there is no data inaccuracy. Problems ensue, however, because there is failure to consider a tornado warning. (A failure of conditional assessment). A third possibility is that SECURE makes no error in data collection nor conditional assessment, yet still goes wrong in interpretation. They might judge that 7 fire alarms is appropriate when 15 are needed. To do their job, SECURE must collect good data, carefully apprise conditions, and then use both of the aforementioned to arrive at an all things considered risk assessment. Note that a safe event is not enough to fend off criticism. SECURE'S customers can demand a refund upon discovery the event unknowingly presented a high safety risk, even if no risk

actualized. Each of us, when making an epistemic risk assessment, functions in a manner similar to SECURE.

Compare a college soccer recruiter, Scott, to SECURE. Scott is asked to watch a promising young athlete, John. Scott could receive inaccurate data via his visual percepts. He might watch the wrong player or mistake made goals for missed ones. Then again, he might have accurate sensory data: a clear view of the right player and what he accomplished. Notwithstanding, problems unfold if Scott mistakes the competition for the best team in the league when it is really the worst (Inaccuracy re: conditions). Lastly, Scott may have accurate data while accurately assessing the relevant conditions. He might still, however, fail in interpretation. New to soccer and hired thanks to nepotism, Scott may think scoring too much reflects poorly on one's soccer abilities. Any of these errors lead to *risk insensitive* belief, and Scott is hence likely to communicate unreliable information.

6 Resolving Lackey's Dilemma

6.1 Morris Versus Henry

With the risk sensitive framework just described, we can now distinguish CHICAGO VISITOR and FAKE BARN: *Morris's belief is risk sensitive but Henry's is not*. Because Morris's belief is risk sensitive, Morris earns credit and acquires knowledge while Henry earns and acquires neither.

As described in Sect. 2, Morris behaves with epistemic excellence insofar as he puts to use a reliable means of true belief acquisition and the use of such means is attributable to him. But now we can describe this in more detail, i.e., in accordance with RSC. The particular ability needed for credit, and that used by Morris, is the ability to assess epistemic risk. Epistemic excellence just is this accurate risk assessment. As previously mentioned, Morris's past experience with, or awareness of, the social practice of directional inquiry makes it possible for him to engage in such risk assessment activity. Assuming this an ordinary circumstance and Morris an ordinary fellow, his automated and conscious cognitive processes assess epistemic risk. Before even asking, Morris observes that the passerby is sane, sober, and human. He then gauges that the offered advice sounds reasonable. Familiarity with the practice bolsters his confidence; in one way or another, life has taught Morris those unqualified to answer directional inquires usually admit as much. All things considered, Morris (using his abilities) assesses (accurately enough) that the passerby's directions are unlikely erroneous. Such risk assessment is an act of epistemic excellence which is attributable to Morris. In this way, Morris earns credit and acquires knowledge.

Henry, like Morris, receives data in need of analysis. For Morris the data was testimony and prior testimonial experience. Morris then gauges the meaning of the data in light of the environment and other relevant conditions. Henry goes through a similar process. He receives data from a visual stimulus and his perceptual system gauges epistemic risk. Henry, however, assumes he is in a traditional barn environment; this skews assessment.

The evaluation of epistemic excellence should be thought of in terms of “total risk assessment.” In other words, epistemic excellence is the proper (or proper enough) processing of relevant epistemic information. An agent receives various information, commonly from many sources and often over long periods of time. Some of this information is consciously accessible, much is not. An agent deserves credit (and so acquires knowledge) when she first processes this data with reasonable accuracy, second comes to the (correct) conclusion that not-*p* is improbable, therefore believes *p*, and *p* is true. Henry, however, misinterprets a *critical* portion of epistemic information when he misjudges his environment as normal barn country. Now one need not have complete or wholly accurate to form a justified belief, but it must be complete and accurate enough to result in a reasonable risk assessment. Because Henry’s missing information was so critical, his assessment falls through.

We might be tempted to think that Henry’s belief forming mechanism is nothing more than visual perception, and this would lead us to conclude he forms his belief via reliable means. But things are not so simple. For instance, in challenging Fred Dretske’s argument against closure, Pritchard has pointed out that beliefs ostensibly formed, “just by looking,” are in reality much more complex. Suppose Zula looks at a zebra and forms the true belief that what she sees is a zebra. It may be tempting to say she forms her belief, “just by looking.” But as Pritchard explains, this isn’t quite right.

I think that while there is a sense in which it is obviously true that Zula gains her knowledge just by looking...perceptual knowledge can...involve a wide range of specialist expertise and background knowledge...such expertise and background knowledge would surely have ramifications for the total evidence that you possess in support of your belief... to know a proposition just by looking need not entail that the only evidence you possess for your belief is the evidence you gained from the bare visual scene before you (2010: 256–257).

Like Zula, Henry’s “evidence” (what I prefer to call data or total information) consists in much more than just the bare visual scene before him. Background knowledge plays an important role; only from past experience does Henry know his percept has the appearance of an object called a “barn,” and only from background information can he judge that open grassy areas are the types of places where barns are commonly found. Yet unfortunately for Henry, some of his background information misleads. If we assume Henry an ordinary fellow, he has not any reason to think that objects that appear like barns are actually barn facades. As far as he knows, it would be pointless to have a town full of barn facades, he has never heard of such things, and he would be prone to suspect (quite reasonably) that those who believe in Fake Barn Country are conspiratorial loons. While these are all reasonable assumptions on his part, they have distorting consequences on his epistemic evaluation.

Total risk assessment is derived from various sources of epistemic information which are first individually interpreted and then collectively assessed. Going too far off the mark when interpreting information will corrupt the collective assessment. This is what happens with Henry. He misinterpreted his environment and this

misinterpretation played a key role in his total risk assessment. Epistemic excellence does not allow for these types of severe mistakes. At least, the type of epistemic excellence that I am singling out does not. In line with previous credit theorists' emphasis on "credit for *success*," an understandable epistemic mishap is still a mishap. The idea is similar to the common externalist/reliabilist notion that justification goes beyond that which is internal to the believer. Even if an agent has good reason to think her method is reliable, she cannot be justified if it's unreliable. Similarly, even if we can understand why Henry made the risk assessment that he did, it was inaccurate and therefore not excellent.

Now some might wonder if the theory I am proposing is really a credit theory, for it seems unfair to blame Henry for inaccurately assessing risk. After all, if Henry is a regular guy and we assume he has a history much like our own history, there is no reason for Henry to think he is in Fake Barn Country. Because of this, there is no reason to think he did anything "epistemically wrong" when he miscalculated the epistemic risk. Said differently, it seems that Henry, while perhaps not justified, is at least warranted. And if Henry did nothing wrong, why should he be denied credit?

This is a good opportunity to clarify the type of credit at work in my theory. The credit relevant to my theory is credit for the excellent performance of a skill (the skill of risk assessment). To stick with the theme of this paper and others in the credit literature, I will use a sports analogy to explain. Imagine an Olympic long jumper who is by most accounts the best in the world. During his jump, against all odds, there is the sound of an elephant in the distance. This startles the jumper and because of this he starts his jump over the line and is disqualified.

The comparison I am trying to draw is one of analogy, where both Henry and my athlete could be understood as "failing." Henry fails, because he lacks a true belief. My athlete fails because he does not win the race. Therefore, both Henry and my athlete fail, even though there is a sense in which they both did everything right. So what Henry and my athlete have in common is the following:

1. They both "failed." Henry did not acquire knowledge (at least, he didn't according to one popular interpretation that I am working with) and the athlete did not win the race.
2. Both of them failed via no fault of their own. Rather, each failed via bad luck that they could not control.
3. Because their failures were due to bad luck, these failures say nothing about each of their given skills. That is, my athlete's failure says nothing negative about his skill as an athlete. Similarly, Henry's failure says nothing negative about his skill as an epistemic agent.
4. Now despite of 3, there is a different sense in which both of their failures knock them down a notch in their relevant field. That is, because my athlete did not win the race, his world ranking as an athlete goes down. He simply is not the world champion. Because winning championships is part of what it means to be a great athlete, there is a sense (despite 3) that he is not as great. Likewise, if there was a world ranking for epistemic agents, Henry would go down a few notches. What it means to be a successful epistemic agent includes acquiring

knowledge, and in this instance Henry just didn't succeed. (even though he has a true belief).

So my point is that despite of competent performance during normal conditions, both Henry and my athlete fail. Both of them fail due to bad luck. Henry fails to know because his risk assessment is off, even though this is due to no fault of his own. His risk assessment is still not taking account of the high risk of fake barns, even though this is nothing he could be blamed for. Similarly, the reason the athlete lost the race is nothing blameworthy.

6.2 Clarifications

Let me make clear that RSC is not a variant of the so called “no false lemmas” theory. As some may recall, shortly after Gettier introduced his problem, a view often referred to as the “no false lemmas” approach (NFL) suggested a simple solution.¹⁴ According to NFL, Gettier's examples of troublesome beliefs are, in actuality, illegitimate (or unjustified) because they rely on false premises: Smith's true belief that “the man who will get the job has 10 coins in his pocket” is acquired by reasoning through the false premise that “Jones will get the job.” Similarly, Smith's true belief that “Either Jones owns a ford or Brown is in Barcelona”, is acquired only via reasoning through the false premise that “Brown is in Barcelona.” NFL proponents argued that a necessary condition of knowledge was that the “belief” in “justified true belief” could not be acquired by reasoning through false premises. With this requirement, we see that the heroes of Gettier's puzzles depend on false premises and therefore their beliefs are not knowledge.

Many problems with NFL soon came to light. First, with some imaginative effort, it is possible to come up with examples similar to those in Gettier's original paper that *do not* rely on false premises.¹⁵ And second, a new breed of cases (sometimes called Gettier cases, although this not all philosophers endorse this interpretation), those of the fake barn variety, were introduced onto the epistemological stage.¹⁶ It seemed to many that simple visual beliefs (like the barn façade belief) do not rely on any premises at all, and hence even more so do not rely on false premises.

Because I emphasize the role false information plays in skewing risk assessment, some might confuse RSC with NFL. I want to be clear that RSC is entirely distinct from, and bears very little relation to any variant of the no false premise approach and *does not* suggest that NFL is necessary for knowledge. Let us return to Henry. I argued that his true barn belief, which might appear to arise spontaneously, is actually dependent on a vast array of background information, much of which is really misinformation. Such misinformation plays a critical role in tipping Henry's risk assessment scales in the wrong way. However, we should not understand Henry's risk assessment failure in terms of false premises. First off, this would make knowledge requirements unreasonably strict. Much of our everyday knowledge is,

¹⁴ See Armstrong (1974, 152–153).

¹⁵ See Lehrer (1965), and Feldman (1974). For a challenge to Lehrer and Feldman, See Levin (2006).

¹⁶ See Goldman (1976).

in all likelihood, partly based on false or misleading background information. For instance, suppose that my belief that George Washington was the first president of the United States is, in part, based on the false background assumption that my kindergarten teacher was an honest broker. This false assumption dings my risk assessment, but *not enough* to curtail my quest for knowledge. My risk assessment can take the hit: I have enough non-misleading information about George Washington that my overall assessment maintains the accuracy required for knowledge.

Not only is false background information compatible with knowledge, it is unclear that background information necessarily consists of beliefs (beliefs to potentially serve the role of a false lemma). Our cognitive system can register information that never makes its way into the realm of explicit beliefs, and might not even rise to the level of implicit belief. Nonetheless, background information contributes to assessment of epistemic risk. It is this failure to accurately assess epistemic risk which accounts for Henry's failure to obtain knowledge. To sum things up: in many cases misleading background information (which may or may not consist of false beliefs) is not enough to prevent a reasonable assessment of epistemic risk. In such cases, one might have knowledge partly based on inaccurate information. However, in other instances, (like with Henry) inaccurate information *does* interfere with a reasonable risk assessment, and thus does prevent one from attaining knowledge.

To those who are sympathetic with having abilities relative to an environment, there is a way to frame things (using my theory) from this perspective. From the Risk Sensitive Credit perspective, the environment is playing a unique role in interfering with risk assessment and hence competence. So, I would agree that in one sense Henry does not have the ability in the environment. Yet I would put that in the following terms: Given the environment that Henry is in, he is not able to accurately assess epistemic risk, and because of this he cannot acquire knowledge. Insofar as ability demands accurate risk assessment (and if I was defining ability it would), then Henry does not have the relevant ability.

Let us return to our analogy of the risk assessment company. Imagine that SECURE concludes that there is minimal safety risk at the mansion, but only because the company is unaware of the man-eating grizzly bears that the malicious neighbor has hidden in the basement. Even though a security company could not have anticipated such a bizarre scenario (and so in this sense are blameless), any risk evaluation made without awareness of this environmental feature will interfere with a successful assessment. Similarly, Henry's ignorance of Fake Barn Country interferes with his epistemic risk assessment and hence creditworthiness. Risk sensitivity demands reasonable accuracy regarding data, environment, and other relevant conditions. Mistakes about any of these can result in an assessment that either (1) misrepresents epistemic risk, or (2) makes an accurate assessment but only by luck. Both (1) and (2) are incompatible with creditworthiness and thereby knowledge. In the former case inaccuracy is the problem; in the latter accuracy is powerless because it does not derive from the agent's abilities. Henry's problem is with (1). His mistaken environmental assumptions give rise to an inaccurate assessment and he gravely misrepresents epistemic risk.

7 Conclusion

Lackey's accusation that credit theory is either too strong or too weak has been hard to overcome. This paper argued that with a few clarifications and theoretical adjustments, credit theory can finally defeat her criticism. The first clarification is between credit for effort and credit for excellence. When arguing that credit theory might be too weak, Lackey assumes the former but credit theory should be understood as the latter. This solves the first half of Lackey's dilemma. To put all worries to rest, credit must be understood in terms of risk sensitivity. We can then see that Morris deserves credit because his belief is risk sensitive while Henry's is not. Maybe Henry is worthy of a certain type of credit, perhaps, we can credit him with an "A for effort." Notwithstanding, because knowledge requires an "A for excellence," effort was just not enough.

References

- Armstrong, D. (1974). *Belief, truth and knowledge*. Cambridge: Cambridge University Press.
- Butts, E. (2014). Slim is in: An argument for a narrow conception of abilities in epistemology. *Journal of Philosophical Research*, 39, 51–66.
- Feldman, R. (1974). An alleged defect in Gettier counter-examples. *Australasian Journal of Philosophy*, 52(1), 68–69.
- Geisler, W. S., & Kersten, D. (2002). Illusions, perception and Bayes. *Nature Neuroscience*, 5(6), 508–510.
- Goldman, A. I. (1976). Discrimination and perceptual knowledge. *The Journal of Philosophy*, 73(20), 771–791.
- Greco, J. (2009). The value problem. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic value*. Oxford: Oxford University Press.
- Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge: Cambridge University Press.
- Greco, J. (2012). A (different) virtue epistemology. *Philosophy and Phenomenological Research*, 85(1), 1–26.
- Henderson, D., & Horgan, T. (2013). Risk sensitive animal knowledge. *Philosophical Studies*, 166(3), 599–608.
- Lackey, J. (2007). Why we don't deserve credit for everything we know. *Synthese*, 158(3), 345–361.
- Lackey, J. (2009). Knowledge and credit. *Philosophical Studies*, 142(1), 27–42.
- Lehrer, K. (1965). Knowledge, truth and evidence. *Analysis*, 25, 170.
- Levin, M. (2006). Gettier cases without false lemmas? *Erkenntnis*, 64(3), 381–392.
- Pritchard, D. (2010). Relevant alternatives, perceptual knowledge and discrimination. *Nous*, 44(2), 245–268.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109, 247–279.
- Riggs, W. (2009). Two problems of easy credit. *Synthese*, 169(1), 201–216.
- Sosa, E. (2009). Knowing full well: The normativity of beliefs as performances. *Philosophical Studies*, 142(1), 5–15.
- Stansbury, D. E., Naselaris, T., & Gallant, J. L. (2013). Natural scene statistics account for the representation of scene categories in human visual cortex. *Neuron*, 79(5), 1025–1034.