

# Fallibilism and the value of knowledge

Michael Hannon

Received: 2 April 2013 / Accepted: 10 July 2013 / Published online: 23 July 2013  
© Springer Science+Business Media Dordrecht 2013

**Abstract** This paper defends the epistemological doctrine of fallibilism from recent objections. In “The Myth of Knowledge” Laurence BonJour argues that we should reject fallibilism for two main reasons: first, there is no adequate way to specify what level of justification is required for fallible knowledge; second, we cannot explain why any level of justification that is less than fully conclusive should have the significance that makes knowledge valuable. I will reply to these challenges in a way that allows me to make progress on a number of important issues in contemporary epistemology: epistemic value, the functional roles of knowledge attributions, experimental epistemology, skepticism, the Gettier problem, and the lottery paradox. My argument is motivated by appealing to various insights derived from the method of ‘practical explication’, particularly the idea that a central purpose of the concept of knowledge is to flag reliable informants. My conclusion is that various practical and theoretical considerations derived from the method of practical explication support the fallibilist conception of knowledge.

**Keywords** Fallibilism · Knowledge · Value · True belief · Gettier · BonJour

## 1 Introduction

A prevailing trend in epistemology over the past several decades has been the acceptance of fallibilism about knowledge. ‘Fallibilism’ is the view that one can have knowledge that a particular claim is true even though one’s justification (evidence, warrant, or supporting grounds) for that claim is less than conclusive (BonJour 2010, p. 57). Put differently, the level of justification requisite for knowing that  $p$  is compatible with

---

M. Hannon (✉)  
Faculty of Philosophy, University of Cambridge, Cambridge CB2 1ST, UK  
e-mail: mhannon09@gmail.com

$p$ 's being false. Most epistemologists have endorsed fallibilism to avoid the skeptical implications of the Cartesian conception of knowledge, which requires the highest possible degree of justification in order to know.

A more recent trend in epistemology is the view that considerations about epistemic value have important ramifications for investigations into the nature of knowledge. A growing number of epistemologists have argued that an essential condition for any adequate theory of knowledge is to make sense of the supposed value of knowledge.<sup>1</sup>

In “The Myth of Knowledge”, Laurence Bonjour appeals to value-related considerations to argue that “this whole ‘fallibilist’ view of knowledge is mistaken” (*ibid*). His main thesis is as follows:

There simply is no well-defined, intellectually significant concept of knowledge fitting this general conception: none that can be genuinely found in commonsense or indeed can be constructed or stipulated in a satisfactory way. The supposed weak [fallible] concept of knowledge in question is, I am suggesting, a philosophical myth. (*ibid*)

Why does Bonjour think the fallible concept of knowledge is a myth? His argument draws a distinction between a ‘conception’ and a ‘concept’: a conception is a general characterization that one or more specific concepts might fit. The fallible conception of knowledge is perfectly real, according to Bonjour, but a specific concept of fallible knowledge would have to say more precisely just what level of non-conclusive justification is required for knowledge. Bonjour thinks that no specific concept can satisfy the fallible conception of knowledge for two reasons: first, we cannot specify what level of justification is required for fallible knowledge; second, and more significantly, even if we could identify the precise level of justification needed for knowledge, we could not say why any such level would have the value or significance that we intuitively take knowledge to have.<sup>2</sup> Thus, such a concept is a myth.

This paper will evaluate and criticize Bonjour’s rejection of fallibilism. Section 2 outlines the two main challenges that Bonjour raises against fallibilism, as well as some of the alleged benefits of rejecting fallibilism. Sections 3–9 reply to these challenges and show that the perceived benefits of rejecting fallibilism are illusory. My aim is not only to provide a reasonable line of response to Bonjour’s challenges, but also to make headway on a variety of important issues in contemporary epistemology, including the functional roles of knowledge ascriptions, epistemic value, experimental epistemology, skepticism, the Gettier problem, and the lottery paradox.

## 2 Two challenges for fallibilism

Let’s grant for now that knowledge is something like (non-accidental) justified true belief, which is an account that Bonjour considers “broadly correct, albeit obviously in

<sup>1</sup> For a discussion of this topic, see Haddock et al. (2010) *Epistemic Value*.

<sup>2</sup> Pritchard (2010, p. 8) calls this the ‘tertiary value problem’. This problem arises because the difference in value between knowing and lesser epistemic states is not thought to be a matter of degree but rather of kind.

need of much additional clarification and specification” (*ibid*, p. 58). With this sketch of knowledge in place, BonJour outlines four traditional assumptions about knowledge (*ibid*, pp. 58–59). First, knowledge is a “supremely valuable and desirable cognitive state” indicative of “full cognitive success”. A satisfactory concept of knowledge must therefore make sense of this supposed supreme value. Second, knowledge is not a matter of degree: either one has it or one does not, with no room for gradations. Third, epistemic justification is a matter of degree. Fourth, epistemic justification is tied to likelihood or probability of truth, with stronger degrees of justification making the truth of the relevant belief correspondingly more likely or probable.

It is not obvious that all of these assumptions are correct. For example, several philosophers have recently argued that knowledge is not the “epistemic *summum bonum*”, as BonJour calls it. Rather, it has been suggested that *understanding*, not knowledge, is the most desirable and distinctively valuable epistemic state (Kvanvig 2003; Pritchard 2009).<sup>3</sup> Also, some philosophers have defended the idea that knowledge *does* come in degrees (e.g., Hetherington 2001). But let’s grant all of BonJour’s assumptions for the sake of argument.

With these assumptions in place, BonJour distinguishes two conceptions of knowledge: the Cartesian (or “strong”) conception and the fallible (or “weak”) conception. More specifically,

*The Cartesian Conception* Knowledge requires the highest possible degree of justification; viz., justification that is *conclusive*, which guarantees the truth of the claim that is believed.

*The Fallible Conception* Knowledge requires a degree of justification that falls short of conclusive justification but is still fairly strong; viz., justification that makes the belief very likely to be true, but does not guarantee it.

The standard for knowledge required by the Cartesian conception is impossible to satisfy in the majority of circumstances (bracketing simple a priori claims and some aspects of immediate experience). Consequently, the Cartesian conception entails that we have very little knowledge. This consequence is often thought to be an unacceptable form of skepticism, which has led many epistemologists to adopt the fallible conception of knowledge. Thus, according to Siegel (1997, p. 164), “we are all fallibilists now”.

But the fallible conception of knowledge is not without problems. This view is committed to the claim that there is some specific level of justification that is less than conclusive but that nonetheless alters our cognitive situation in an important way. Imagine that we are inquiring as to whether *p* is true. It is easy to understand how finding increasingly high levels of justification for the claim that *p* improves our cognitive situation by making it more likely that our belief is true; but the idea that there is some specific level at which our cognitive situation transforms from not-knowing to knowing seems more than a little peculiar. BonJour calls this the “magic” level of justification. Before this level is attained one merely has a belief that is more

<sup>3</sup> The existence of trivial truths also suggests that knowledge is not always more valuable than true belief, since it is not obviously better to know trivial truths (Sosa 2003).

and more likely to be true, but at the “magic” level one suddenly and miraculously has knowledge.

We can put this point differently by thinking of probability as measured by the use of decimals in the interval  $[0, 1]$  on the number line. A probability of 0 means that the claim is guaranteed to be false and a probability of 1 means that the claim is guaranteed to be true. Framed this way, BonJour’s (2010, p. 61) question is: how probable must your belief be to qualify as knowledge? He thinks “it is fair to say that nothing like a precise specification of the ‘magic’ level has ever been seriously suggested, let alone more widely accepted”. Providing such a precise specification has at least two difficulties: first, any point lower than 1 seems arbitrary (why pick *that* point precisely?); second, it is unclear why achieving this specific level would make such a difference in our cognitive situation (by taking us from not-knowing to knowing).

BonJour admits that it might be “unreasonable to demand a precise numerical specification of a level of probability and perhaps not even entirely clear that numerical probabilities are the right way to think of degrees of justification” (*ibid*, p. 60). But presuming that there is a definite concept of knowledge that denotes an “exalted cognitive state”,

it is surely not good enough to say merely, as is commonly said, that the level of justification in question is “strong” or “high” or “adequate” or enough to make it “highly likely” that the belief in question is true, for nothing this vague is enough to specify a definite level of justification and a corresponding definite concept of knowledge. (*ibid*)

We can summarise BonJour’s two challenges in the following way:

*The Conceptual Challenge* What specific degree of justification is required to achieve the exalted cognitive state of knowledge?

*The Value Challenge* Why does achieving this specific level of justification make such a difference and why is this difference so important?

BonJour considers the value challenge to be more serious because he thinks it is “impossible to see what could give any level of justification that is short of being conclusive the kind of special significance that the weak [fallible] conception requires it to have” (*ibid*, p. 61).

We are therefore led to reject fallibilism by reflecting on the value of knowledge. On the fallibilist picture, it is clear that further increases in justification do not cease to matter even once knowledge has been obtained. Presuming that the level of justification is less than conclusive, additional increases in justification will still be cognitively valuable *in exactly the same way* as earlier increases in justification, before the “magic” level was reached (*ibid*). So why is it so important whether or not the “magic” level is reached? Why is this level of justification special? BonJour thinks there is no plausible answer to this question, no account of why any non-conclusive level of justification would have such a special status.

Furthermore, BonJour argues that the fallible conception of knowledge gives rise to two additional problems: the ‘Gettier problem’ and the ‘lottery paradox’. If we abandon the fallible conception, however, these problems allegedly disappear.

Gettier (1963) showed that there are cases in which an agent’s true belief is strongly justified and yet intuitively it does not qualify as knowledge. These so-called ‘Gettier cases’ arise because the level of justification required for knowledge is presumed to be less than conclusive, which allows luck to undermine knowledge. For example, Smith might have a strongly justified true belief that the man who will get the job has ten coins in his pocket, yet Smith’s belief is true in a way that intuitively is merely accidental in relation to his justification.<sup>4</sup> If knowledge requires conclusive justification, however, there would be no room for luck and thus Gettier cases would pose no problem.<sup>5</sup>

Moreover, even if we found some fourth condition to make knowledge Gettier-proof (e.g., non-defeasible justified true belief<sup>6</sup>), the addition of such a condition would not yield any insight into *why* knowledge should require such a condition. What would the rationale for such a condition be, apart from merely avoiding counterexamples? Without any rationale for why knowledge should require the elusive anti-Gettier condition, the significance of the concept of knowledge remains obscure. We can avoid these problems by denying the assumption that knowledge requires less than conclusive justification.

BonJour also targets fallibilism as an essential ingredient to the so-called ‘lottery paradox’ (see Hawthorne 2004, pp. 3–7). Most people agree that any person who believes on the basis of high probability that he or she will lose a fair lottery does not qualify as knowing that he or she will lose. But ruling out knowledge in the lottery case seems to impugn a much wider range of ordinary, commonsense knowledge.<sup>7</sup> For example, it seems acceptable for me to claim to know that (*p*) I will be in Cambridge this summer, despite the fact I do not know that (*q*) I will not die before then. Yet (*p*) entails (*q*); so, given closure, I cannot know (*p*) without knowing (*q*). However, I do not know (*q*); therefore, I do not know (*p*).

Can fallibilists solve the lottery problem? BonJour (2010, p. 68) says that the “deepest and most serious objection” to any proposed additional requirement to rule out

<sup>4</sup> Here is a fuller description of the original case from Gettier (1963). Suppose that Smith and Jones have applied for a certain job, and suppose that Smith has strong evidence for the following proposition: (a) Jones is the man who will get the job and Jones has ten coins in his pocket. Smith’s evidence for (a) is that he counted the coins in Jones’s pocket 10 minutes ago and that the company president assured him, Smith, that Jones would be selected. Proposition (a) entails (b): the man who will get the job has ten coins in his pocket. Let us suppose that Smith sees the entailment from (a) to (b) and accepts (b) on the grounds of (a), for which he has strong evidence. In this case, Smith is clearly justified in believing that (b) is true. But imagine, further, that unknown to Smith, he himself, not Jones, will get the job. Also, unknown to Smith, he himself has ten coins in his pocket. Smith has a justified true belief that does not seem to be knowledge.

<sup>5</sup> Following Zagzebski (1994, p. 66), Pritchard (2005, p. 149), and Turri (2011, p. 1), I take Gettier cases to involve a ‘double luck’ structure. We start with a belief that meets the justification requirement for knowledge; then add an element of bad luck that would normally prevent one’s justified belief from being true; finally, add an element of good luck that cancels out the bad luck, so the belief ends up true.

<sup>6</sup> The justification for a proposition *p* is defeasible if there is some true proposition *q* such that if the person were to become justified in believing *q*, he or she would no longer be justified (to the requisite degree) in believing *p* (BonJour 2010, p. 65).

<sup>7</sup> Vogel (1990) highlights this point.

lottery cases would be the lack of any plausible intuitive rationale for this requirement. We have no clear understanding of why knowledge should involve such a requirement. If we endorse infallibilism, however, then the lottery example constitutes a simple and decisive *reductio* against fallibilism. If less than conclusive justification were enough for knowledge, then we would have knowledge in the lottery case; but we intuitively do not. Why, then, do we have the intuition that we possess knowledge in other (i.e. non-lottery) cases? BonJour says it is because “this lack of conclusiveness is less conspicuous and so more easily ignored” (*ibid.*, p. 70).

BonJour therefore recommends that we abandon the fallibilist conception of knowledge. He then tentatively argues that the correct conception of knowledge is some version of the Cartesian view, which holds that one knows that  $p$  only if one’s justification guarantees that  $p$ . This view is said to have several advantages over fallibilism. In addition to avoiding the Gettier and lottery problems, we also get straightforward answers to BonJour’s two challenges. We have a clear answer to his conceptual challenge, which asks us to specify the degree of justification required for knowledge: we must have conclusive justification. Moreover, we have a plausible answer to his value challenge, which asks why achieving knowledge-level justification is so intellectually important and valuable: presupposing that truth is the aim of belief, it makes sense to say that knowledge requires conclusive reasons because they would guarantee that the belief’s aim is achieved. Complete cognitive success is achieved only when our justification is sufficient to establish completely that our belief is true, which makes it natural to identify the situation of such completeness as the most valuable cognitive state, namely knowledge (*ibid.*, p. 59).

The challenges for fallibilists are (a) to offer a clearer specification of the “magic” level of justification, (b) to explain why this level has its special status, and (c) to show that the perceived benefits of infallibilism vis-à-vis the Gettier and lottery problems are illusory. I will develop a version of fallibilism that answers these objections—a task that BonJour considers “impossible” (*ibid.*, p. 76). More specifically, I will respond to BonJour’s first challenge by arguing that commonsense tacitly involves a concept of fallible knowledge that accounts for the “magic” level of justification to a reasonable degree of approximation; however, I will also suggest that too much precision here should not be expected. I will then explain why the “magic” level of justification has the sort of significance that fallibilism attributes to it. I shall motivate my argument by appealing to the hypothesis that a central purpose of the concept of knowledge is to flag reliable informants (Craig 1990). Finally, I will show that the perceived benefits of rejecting fallibilism are in fact illusory. My conclusion is that there is a non-conclusive level of justification that can play the role required by the fallibilist conception of knowledge.

### 3 A problem of vagueness?

An initial worry is that BonJour’s critique of fallibilism is merely rehearsing the familiar problem of vagueness. Take the classic example of the predicate ‘is bald’. Clearly someone with 10,000 hairs on his head is not bald and clearly someone with no hairs on his head is bald; but at what point does a person go from being ‘not-bald’

to ‘bald’? Must there be less than 100 hairs on his head? Less than 99 hairs? 98? etc. It is reasonable to think that there is no precise point at which someone goes from being not-bald to bald: the predicate ‘bald’ is vague.<sup>8</sup> And since this form of sorites-style reasoning leads to absurd conclusions (e.g. there are no bald people or, *mutatis mutandis*, everyone is bald), any such argument is assumed to be problematic.<sup>9</sup>

The analog with justification in terms of probabilities is something like this:

- (1) A probability of 0.01 is not sufficient for knowledge-level justification.
- (2) If a probability of 0.01 is not sufficient for knowledge-level justification, then a probability of 0.02 is not sufficient for knowledge-level justification. ...
- (3) If a probability of 0.98 is not sufficient for knowledge-level justification, then a probability of 0.99 is not sufficient for knowledge-level justification.

If we presuppose that an increase in probability of 0.01 can never make the difference between not-knowing and knowing, then this reasoning will illustrate that there is no such thing as knowledge, since someone with conclusive justification would still lack knowledge. This is a conclusion everyone should reject. (Even the skeptic thinks that conclusive justification is necessary for knowledge; she simply denies that humans can attain this level of justification.) However, if we suppose that a difference in probability of 0.01 can take one from not-knowing to knowing only when it makes the difference between non-conclusive and conclusive justification (i.e. 0.99–1), then the Cartesian view is correct. Either way, fallibilism is false.<sup>10</sup>

BonJour might argue that his main objection to fallibilism has nothing to do with vagueness because we may assume that there *is* a precise point at which the appropriate level of non-conclusive justification transforms one’s true belief into knowledge. Presume, for instance, that a probability of 0.90 is the cut-off for knowledge-level justification. BonJour’s main objection to fallibilism is that, even if there were such a precise boundary, it would be *arbitrary* for this “magic” level to be the point that separates not-knowing from knowing. It would be arbitrary because minor increases in justification, whether before or after this threshold, are cognitively valuable in precisely the same way; thus, we cannot explain why knowledge-level justification has special significance. For this reason he might claim that the issue is one of arbitrariness, not vagueness.

However, this reply does not show that BonJour’s challenge differs from familiar vagueness worries. For example, ‘personhood’ is not a precise notion with well-defined boundaries. But even if we assume that there is a precise point at which a fetus becomes a person (i.e. 100 days), this “magic” point would still seem arbitrary. Thus, BonJour does not appear to be offering an objection that is any different in structure from familiar vagueness-related objections that have been raised in other areas.

<sup>8</sup> Williamson (1994) is an exception.

<sup>9</sup> If sorites arguments succeed, the loss of fallibilism would be the least of our worries. We would have problems to deal with for virtually any concept that admits of degrees (or any concept that has a necessary component that admit degrees, as with the justification condition for knowledge).

<sup>10</sup> Modeling justification by the use of decimals on the number line is merely a helpful idealization. We should not interpret BonJour as demanding a degree of precision stated in such terms. Rather, his point is that a *small* increase in justification cannot take one from not-knowing to knowing in a non-arbitrary way that accounts for the value of knowledge, unless it leads us to conclusive justification.



Nevertheless, fallibilists must still answer BonJour's challenge. It is surely not good enough to say that the level of non-conclusive justification required for knowledge is "high" or "strong" or makes one "highly likely" to be right. More needs to be said if we are to have a cognitively significant concept of knowledge that fits the fallible conception. Thus, I take BonJour's challenge seriously.

#### 4 A reasonable degree of precision

Does BonJour demand more than a reasonable degree of precision? I think he does. To illustrate why, consider the concept of an 'expert'. We might ask, what is the relevant standard to qualify as an 'expert' craftsman? How many faultless constructions must she make? How many different objects must she know how to craft? In other words, what is the "magic" level of skill required for expertise in this case? There seems to be no clear answer to this question. There is no obvious and seemingly non-arbitrary point at which one who exercises a skill becomes an expert at it.<sup>11</sup> Does this suggest that there are no experts? Of course not. Does it suggest that in order to qualify as an expert one must be *maximally* skilled such that one could not possibly be more skilled? That also seems false (and perhaps incoherent). Why, then, should we expect knowledge to be any different?

The comparison between an expert and a knower can be strengthened. Notice that 'expert' and 'knower' are both evaluative concepts that serve important roles in our socio-linguistic interactions. More specifically, both concepts are used to flag people who exhibit a particularly high degree of reliability in some domain. An expert craftsman successfully creates the products of her craft in a reliable manner due to her abilities (or at least she could in the appropriate conditions). Similarly, a knower is someone who successfully reaches true belief in a reliable manner due to his cognitive abilities. The comparison between expertise and knowledge is therefore quite apt.<sup>12</sup> (Plato's term *techné* is often translated as "knowledge", but [Gentzler 1995](#) and [LaBarge 1997](#) suggest that it is perhaps better translated as "expertise".)

This raises a challenge to BonJour's view. His objections to fallible knowledge seem to apply equally well to other concepts, such as the concept of an expert. So unless we think it is plausible to accept infallibilism about expertise—i.e. the view that in order to count as an expert one must be maximally skilled such that one never makes any mistakes and could not improve whatsoever—BonJour's argument looks less convincing. It seems logically possible for experts to exhibit even more skill than they do. This puts some pressure on BonJour's anti-fallibilist argument.

<sup>11</sup> One might not find the expertise metaphor convincing because there might be clear, regimented, and institutionalized answers to this question in some trades. This point does not obviously apply to my case of the craftsman, and I suspect there are many areas in which there is no clear standard or test for expertise. Moreover, I worry that any institutionalized answer would seem arbitrary. Does passing a test about the physical and chemical properties of propane suffice to make one an expert even if one cheated on the test or luckily guessed many of the answers? In any case, nothing significant turns on my expertise example. Similar worries apply to who counts as a 'philosopher', for instance. There are many other such concepts, such as the concept of an 'addict'.

<sup>12</sup> This analogy is not perfect if we assume that knowledge does not come in degrees, since it seems natural to say that A can be more of an expert than B.



So where does BonJour's argument go wrong? The aforementioned comparison between knowers and experts suggests that he demands too much precision from our evaluative concepts. How much is too much? This is a difficult question to which I can provide no pithy answer (although I will develop my answer to this question in Sects. 5 and 6). BonJour (2010, p. 70) thinks "there is no apparent way to specify or determine, even to a reasonable measure of approximation, what the required non-conclusive level of epistemic justification might be". I doubt that we can specify the level of skill needed to qualify as an expert in some domain any more precisely than we can determine the required level of justification for knowledge. This suggests that BonJour mistakenly expects an analysis of such concepts to provide more precision than is achievable (or necessary).

What, then, is a "reasonable measure of approximation" for specifying the non-conclusive level of justification needed for knowledge? To answer this question, we must first reject a presupposition of BonJour's argument. He thinks of justification in terms of gradual increase, which presupposes a continuous *quantitative* model of epistemic justification. This view renders mysterious why the particular point that knowledge marks out is more significant than a point just before or after it. BonJour therefore concludes that there is no adequate answer to his challenge, since it is implausible to maintain that a gradual increase in justification could yield a qualitatively different state, namely knowledge.

However, we can reject the quantitative model presupposed by BonJour and instead adopt a view that understands justification *qualitatively*, with potentially significant differences between adjacent levels of justification. If knowledge is qualitatively different from lesser epistemic states, then it makes sense to think that justification should be understood qualitatively rather than quantitatively. By conceiving of justification in these terms, we can reply to BonJour's challenges.

## 5 The qualitative model of justification

Conee and Feldman (2004) have endorsed the qualitative model of justification in their book *Evidentialism*. They recommend the legal standard for conviction in criminal cases as sufficient for knowledge-level justification, which is proof beyond a reasonable doubt (2004, p. 296). This level of justification is weaker than, say, a mathematical proof, but it is stronger than a good reason to believe; it calls for "a strong reason to believe". For a belief to meet this standard, one must have "strong reasons in support of it, no undefeated epistemic reasons to doubt it, and no undefeated epistemic reason to believe that one's evidence for it is unreliable" (*ibid*). Two features of this account are important for my purposes: first, Conee and Feldman do not require conclusive justification and thus defend a version of fallibilism; second, this account leaves open exactly how strong one's reasons must be in order to qualify as strong enough. As they say, "[t]here is no conspicuously correct fact of this matter" (*ibid*).<sup>13</sup>

<sup>13</sup> Non-evidentialists can also endorse a non-quantitative model of justification. I simply mention evidentialism for illustrative purposes.

Conee and Feldman provide a clear statement of the framework of epistemic justification within which I shall operate. They do not conceive of epistemic justification in terms of gradual increase but rather in terms of qualitatively discrete categories: ‘good reason to believe’, ‘strong reason to believe’, ‘deductive certainty’, etc.<sup>14</sup> According to Conee and Feldman, there is one epistemic position in which an agent must stand with respect to a proposition in order for that agent to qualify as knowing it. Contextualists will deny this. According to contextualism, whether an agent qualifies as knowing will vary with features of the conversational context. Thus, while the legal standard of conviction in criminal cases might be appropriate in the context of a criminal trial, it is not the operative standard in all contexts.

Although I favor contextualism, my aim in this paper is not to argue for this view. Rather, I will show that both invariantists and contextualists can endorse the qualitative model of justification. For ease of exposition, I will temporarily bracket contextualism when discussing the qualitative model of justification (and I’ll return to it soon). In what follows, I will specify the non-conclusive level of justification that is required for knowledge and also explain why this level can account for the value or importance that we attribute to knowledge. I will motivate this view by appealing to the increasingly popular idea that the concept of knowledge is important because it allows us to identify a certain type of reliable informant.

## 6 The reliable informant standard for knowledge

I propose the following non-conclusive level of justification as necessary for knowledge:

### *The Reliable Informant Standard for Knowledge*

The level of justification needed for knowledge is that which puts the agent in a strong enough epistemic position for her to fittingly serve as a reliable source of actionable information for members of her epistemic community, many of whom have diverse projects, purposes, and interests.<sup>15</sup>

This view is motivated by Craig’s 1990 book *Knowledge and the State of Nature*. In that work, Craig hypothesizes that the concept of knowledge is valuable because it enables us to flag reliable informants (to put his central thesis briefly and roughly). This hypothesis derives from the idea that the concept of knowledge functions to facilitate the survival and flourishing of human beings in communities. Humans need true beliefs about their environment in order to successfully guide their actions, so they require sources of information that will lead to true beliefs. Often the most efficient way to acquire a true belief as to whether  $p$  is to ask someone reliable whether  $p$ ; and since on any issue some informants will be more likely than others to provide a true

<sup>14</sup> Chisholm (1977, pp. 9–10) mentions the following grades of justification for a belief: acceptable, reasonable, beyond reasonable doubt, evident, certain.

<sup>15</sup> You might worry that this just pushes the question of the level of justification required for knowledge to the question of the level of reliability required for someone to fittingly serve as a source of actionable information for members of her community. However, I will answer this worry by specifying that level of reliability.

belief, any community may be presumed to have an interest in evaluating sources of information. Having a word to flag reliable informants is therefore an important role in the general economy of our concepts.

This view is growing in popularity, and a number of epistemologists have tried to derive significant epistemological payoff by endorsing it. For example, Greco (2010) uses the idea that knowledge ascriptions serve to flag reliable informants to defend his view that ‘knowledge is credit for true belief’. Similarly, Pritchard (2010) argues that this hypothesis provides support for his ‘anti-luck virtue epistemology’. In this paper, I will appeal to this hypothesis in order to motivate a version of fallibilism that answers BonJour’s challenges.

It is clear that a ‘knower’ and a ‘reliable informant’ are not synonymous; for example, an informant might qualify as ‘reliable’ for me because I have very little stake in the matter, yet he might not qualify as a knower. Imagine that my brother and I are exploring a city that neither of us has visited in several years. My brother might not *know* whether the nearest coffee shop is down the left or the right side of a fork in the road, but I might regard him as a reliable source of information because his memory for directions is usually quite good. Given this gap between a knower and a reliable informant, how do we derive the concept of knowledge from the concept of a reliable informant?

The concept of knowledge derives from our need to identify reliable informants *in general*. Although each of us needs reliable information to successfully guide our actions, we also have a shared need to pool and transfer information in order to make it easily accessible. We therefore have a need to identify individuals who are reliable not just for me here and now (or for you there and then), but also for a variety of individuals with a wide range of interests, projects, and purposes. There is a need to assess and identify the adequacy of informants for persons and purposes beyond our own immediate interests. However, we often do not know the practical needs or interests of agents to whom an informant may be useful, so we need to ensure that informants are reliable enough for many people and purposes.

Craig (1990, p. 82) appeals to a process he calls *objectivization*. Objectivization occurs when we abstract away from details of the particular circumstances surrounding one individual’s inquiry to arrive at a more intersubjective, socially directed concept. The objectivized notion of a reliable informant is what resembles our familiar concept of knowledge: when we attribute knowledge to someone, we certify that agent as epistemically positioned such that we can freely draw on her information. Knowledge ascriptions certify an informant’s belief or information to a shared epistemic community. As Henderson (2009, p. 126) puts it, “one who certifies sources of information [...] should be concerned with informants with actionable information for an indeterminate range of projects and their various practical interests”. A ‘knower’ must satisfy this general requirement.

The theory of relevant alternatives provides a way to clearly articulate this version of fallibilism. What it takes to qualify as knowing that  $p$  is that the agent is in a strong enough epistemic position with respect to  $p$  to discriminate  $p$ ’s obtaining by eliminating all of the not- $p$  possibilities that are relevant alternatives to members of the epistemic community that might draw on her information. Which alternatives are relevant to the members of our epistemic community and which are not? Many think

that this question has not been answered with a satisfying degree of precision. This difficulty has led critics to view the relevant alternatives theory as ad hoc or obscure. I follow Austin (1946), who thought that we have an idea (rough, to be sure) of what normally counts as having done enough to establish the propriety of a knowledge claim. Certain alternatives are not the kind that we normal humans take to be the likely counter-possibilities to what the subject is said to know, and thus we need not rule them out. For example, it does not seem that I can be legitimately criticized or blamed for failing to rule out the possibility that I am actually a brain in a vat when making my epistemic evaluations, at least not in the same way that I could plausibly be criticized for failing to take into account the possibility that (say) the bird in my garden is a goldfinch rather than a canary.

There is considerable vagueness here. We could try to cash out the reliable informant standard for knowledge by stating, proposition by proposition, and circumstance by circumstance, what would need to be done to establish the propriety of claiming to know any given proposition.<sup>16</sup> Of course this is not a task that we really want to accomplish, nor is it one that we must accomplish. Furthermore, I do not expect there to be unanimous agreement about all of the possibilities that do or do not need to be ruled out. I merely presume that such judgments will (or would) coincide sufficiently to give us what Rysiew (2001, p. 489) calls “a set of ‘core’ not-*p* alternatives”.<sup>17</sup> But there is no reason why relevance cannot be a somewhat vague notion. The point is that we can say tolerably well, in particular cases, what does and what does not need to be eliminated in order to appropriately credit somebody with knowledge.

Is this level of justification specified to a reasonable measure of approximation? I suspect that BonJour will not be satisfied with the level of precision involved in my specification of the non-conclusive level of justification for knowledge. I’ll quote him again to illustrate his dissatisfaction:

it is surely not good enough to say merely, as is commonly said, that the level of justification in question is “strong” or “high” or “adequate” or enough to make it “highly likely” that the belief in question is true, for nothing this vague is enough to specify a definite level of justification and a corresponding definite concept of knowledge. (2010: 60)

This objection applies to some of Craig’s own remarks; for instance, Craig (1990, p. 91) argues that a knower is “someone with a very high degree of reliability, someone who is very likely to be right—for he must be acceptable even to a very demanding inquirer”. However, my view is more precise than the idea that the level of justification is ‘highly likely’. I have appealed to the hypothesis that the concept of knowledge functions to identify reliable informants to firm up the level of justification needed for knowledge. The specified level is admittedly not very precise, but I have questioned whether we should expect more precision when analyzing certain concepts. We are not forced to put a precise value because the subject matter does not demand it. We can specify

<sup>16</sup> This method is suggested (but not endorsed) by Kaplan (2008, p. 353). My position is strongly informed by his view.

<sup>17</sup> Without this assumption it would render utterly mysterious how people are adept at determining what a speaker means when uttering “S knows that *p*” (Rysiew 2001, p. 489).

the required non-conclusive level of epistemic justification to a reasonable measure of approximation when that demand is properly understood.

Here I follow Bertrand Russell's good advice, which cautions us that

'Knowledge' is not a precise conception [...] A very precise definition, therefore, should not be sought, since any such definition must be more or less misleading. (Russell 1912: 134)

In the case of knowledge, people have developed a conception for how reliable an informant must be to count as reliable enough because the practice of sharing information is one that we all grow up into, and which has continued for some time. We manage to become proficient at the practice of distinguishing those possibilities we need to eliminate from those we do not as a result of socialization and acculturation. The practice of knowledge attribution is learned in the same way that humans of sufficient maturity learn to tell when it is appropriate to stop asking "Why?".<sup>18</sup> Inquiry may reasonably terminate when our informant has done enough. How much is enough? Austin reminds us,

Enough is enough: it doesn't mean everything. Enough means enough to show that (within reason, and for present intents and purposes) it 'can't' be anything else, there is no room for an alternative, competing description of it. (1961: 52)

Our everyday thinking about likelihoods does not happen in precise quantitative terms. The workings of our ordinary practices are not neat and tidy enough to provide such a simple, general account that will easily tell us by virtue of what something does (or does not) need to be ruled out in a given case (Kaplan 2008, p. 354).

What is clear is that, despite the imprecision involved, our ordinary practice of knowledge attribution works well in the course of everyday life. Moreover, this account seems highly plausible and a promising way to think about the so-called "magic" level of justification.

If the story I have told is roughly correct, we might ask whether one epistemic standard will govern our knowledge ascriptions in all contexts. In other words, does this view imply that contextualism is mistaken? Several authors have taken Craig's view to support some version of invariantism. For example, Kelp (2011, p. 65) argues that this proposal will lead us to a 'high-standards' version of classical invariantism. Kusch (2011, p. 11) agrees with Kelp that there will be an invariant standard, although he denies that this standard will be particularly high. I think that the reliable informant standard naturally approximates the level envisioned by classical (non-skeptical) invariantism, but nothing I have said rules out contextualism.<sup>19</sup>

<sup>18</sup> This example is due to Kaplan (2008, p. 255).

<sup>19</sup> Henderson (2009, 2011) shows that contextualism is not only compatible with the idea that knowledge attributions flag reliable informants, but also that this idea actually motivates contextualism. He draws a distinction between two broad communities: applied source communities and general source communities (2009, p. 126). The members of an applied source community are focused on some practical project, and are therefore concerned with sources of actionable information on which to proceed in their particular project. The members of a general source community, in contrast, are devoted to developing a body of results on which folk in various other communities might confidently draw, whatever their projects or

The reliable informant standard for knowledge not only specifies the “magic” level needed for fallibilism, it also suggests why this standard is intellectually significant. Inquiry can be thought of as a process of ruling out various possibilities. As inquirers, we seek reliable information on which to base our actions. The basic aim of belief is to store correct information, for the rational person wants his actions to be guided by true beliefs. However, the process of inquiry is potentially open-ended. For example, what does it take to know that the bird in my garden is a goldfinch? Must I see that it has a red face? Must I distinguish it from woodpeckers, which also have red faces? Must I eliminate the possibility that its face is merely painted red? Must I prove that I am not slightly colorblind? Must I know that I am not a brain in a vat? It is *always* possible to continue one’s inquiry. The reliable informant standard provides a general point at which people can reasonably terminate inquiry.<sup>20</sup> This level of justification is cognitively valuable because it satisfies one of the platitudes about the functional roles of knowledge ascription: it signals the point of legitimate inquiry closure (Kvanvig 2009, p. 344). The same cannot be said for belief, justification, reliability, etc.<sup>21</sup>

Kappel (2010, p. 74) argues that Craig does not consider the special role the concept of knowledge has for inquiry independently of whatever need we might have for identifying sources of information. In particular, Kappel argues that a very distinct role for the concept of knowledge in inquiry is our need for an *inquiry-stopper*.<sup>22</sup> Our need for an inquiry-stopper arises out of the following trivial observations about our interest in truth and certain limitations in our cognitive capacities: (1) truth matters; (2) inquiry is needed to acquire truth; (3) inquiry is always costly in the sense that continuing to inquire requires time and resources; and (4) inquiry has no natural stopping point because there are always further conceivable but as yet uneliminated error possibilities. We therefore need a way to command a switch of attention away from these further uneliminated possibilities. We need to signal when inquiry has gone on long enough.

---

Footnote 19 continued

purposes might be. Since the concept of knowledge is used to certify epistemic agents as reliable sources of information to some audience or group, whether an agent qualifies as a knower depends on the audience for which that agent is evaluated as a source of information. This motivates contextualism in the following way: when one is certifying a source of reliable information for an applied source community, one should be concerned with sources of information in view of the practical interests of the applied community. When one is certifying an epistemic agent for a general source community, one should be concerned with sources of information that will be reliable for an indeterminate range of applied communities and their various practical interests (i.e. low or high stakes), not a concrete limited purpose. I mention Henderson’s view to show that contextualism is not obviously ruled out by the position that I am defending.

<sup>20</sup> I say “reasonably” terminate inquiry because I want to rule out cases in which we no longer feel like investigating or because we want to be alleviated from doubt. Also, false belief would not be a successful end to inquiry.

<sup>21</sup> There may be exceptional cases in which speakers are willing to ascribe knowledge to somebody and yet it is reasonable to check further (see Brown 2008, pp. 1444–1445; Reed 2010, pp. 228–229). However, it is uncontroversial that “S knows that *p*” typically conveys that there is no need for further investigation (see Cohen 1999, p. 59; Millar 2010, p. 98).

<sup>22</sup> Kvanvig (2003, p. 171), Kelp (2011), and Rysiew (2012) also argue that knowledge ascriptions terminate inquiry. Kelp and Rysiew argue that Craig misidentifies the main function of knowledge ascriptions by linking the concept of knowledge with identifying reliable informants rather than ending inquiry. I will demonstrate that these two ideas are compatible.

While Kappel is surely right about our need for an inquiry-stopper, my approach explains why the functional role of flagging reliable informants (in the objectivized sense) also serves to mark the point at which further inquiry is unnecessary, as well as why spending more time and resources to continue one's inquiry would be impractical. Continuing to inquire beyond this point would commit us to paying higher "informational costs" that are not worth the lessened risk of being wrong. Everyday life does not demand that our chances of being wrong are absolutely zero. If our informant knows, then there is no need to investigate further. Attributing knowledge to someone is a way of expressing the attitude that someone's epistemic position (with respect to a given proposition) is good enough to stop further inquiry. That's precisely what makes such a person reliable enough.

My view also explains why we should expect practical factors to bear on whether somebody qualifies as a knower, since it is plausible that the higher the stakes are, the more evidence one needs before one closes inquiry. In the context of a murder trial, for example, where the most serious practical consequences turn on being right, our inclination is to heighten the epistemic standard another notch. The epistemic goal of knowledge cannot be understood without taking into account pragmatic features of our lives.

The principle of objectivization leads us to a concept of knowledge that functions to identify informants on whom people can rely for a wide variety of projects and purposes. Will the principle of objectivization push us towards accepting more and more demanding epistemic standards that eventually lead to infallibilism (and into the arms of the skeptic)? If we are looking to identify informants on whom anyone can rely, BonJour might argue that this will lead us to demand infallibility from our informants. After all, an infallible informant is surely reliable enough for any inquirer, even extraordinarily demanding ones. Is this argument plausible?

We are held back from going over the edge into infallibilism because doing so would mark the end of any contact with the practical requirements that comprise the purpose of our concept of knowledge (Craig 1990, p. 117). To insist upon this level of justification would exclude a vast number of inquirers on whom most people could reasonably rely for their purposes. Our withholdings would suggest that these perfectly reliable informants do not have sufficiently good information on which we can base our beliefs and actions. But we rarely require our informants to have the highest quality epistemic position. The practical explication of knowledge is rooted in the socio-epistemological need for pooling and sharing information, which generates a need to flag informants who can provide us with reliable information. Setting the standard too high will frustrate our communal epistemic practices and thus would be anti-thetical to this goal. This answers BonJour's question why we should set the fallible standard at this point and not higher.

Ultimately, whether or not we should reject fallibilism depends on whether there is a better alternative. BonJour regards the Cartesian conception of knowledge as a more plausible view.<sup>23</sup> A full-scale assessment of the merits of fallibilism vis-à-vis

---

<sup>23</sup> Admittedly, defending this position is not the main concern of BonJour's paper. In fact, he is "uncertain" about whether the Cartesian view is correct and he suggests that there may be no coherent concept of knowledge in commonsense (2010, p. 58). Surely, however, an adequate refutation of fallibilism must show



the Cartesian view will not be offered here, but I will demonstrate that on several important issues my view is more plausible than Cartesianism.

## 7 Sophisticated infallibilism

Many philosophers quickly dismiss the Cartesian conception of knowledge because of its skeptical consequences. However, we are not entitled to assume in advance that any theory which has such consequences is guaranteed to be false. Nevertheless, there is a very strong presumption that a theory with radical skeptical consequences does not capture the ordinary concept of knowledge that is actually in use. Even if we suppose it is possible that our concept of knowledge has widely unrecognized skeptical implications, any competing theory that does not have such implications will possess an advantage. This is not to assume that skepticism is false or that we can never reach any conclusion except those that confirm our pretheoretical intuitions; it is simply a criterion of theoretical adequacy. Our non-skeptical intuitions are both pervasive and persistent. Admittedly, these intuitions are unstable to some extent, since skeptical pressure sometimes leads us to withdraw or weaken our knowledge claims; however, it would be misleading to overemphasize these fleeting moments of doubt. Furthermore, fallibilism does not need any extra explanation about why commonsense could be so radically mistaken.

BonJour's defense of infallibilism amounts to what Greco (2000, p. 16) calls "sophisticated skepticism", which tries to explain why we happily (but mistakenly) attribute knowledge in the course of everyday life. Roughly, the idea is that practical purposes make it appropriate to attribute knowledge in a wide range of cases, which misleads us into thinking that such claims are true. Such a view thereby prevents us from denouncing infallibilism on the grounds that this view cannot account for our commonsense intuitions about knowledge.

Assuming for the moment that infallibilism is true, why do people routinely make false knowledge attributions? BonJour (2010, pp. 71–74) attempts to reconcile the Cartesian view with the ordinary usage of 'know' by providing a number of explanations for why what seems obviously true (that we have knowledge) is in fact false. One explanation is that people falsely attribute knowledge due to simple epistemological error: we reasonably regard our justification for a belief as conclusive when deeper philosophical insight shows that it is not. This seems implausible, however, because people often come to terms with their own fallibility. We all recognize that a well-supported belief may yet turn out to be false and that we cannot insulate ourselves from all possibilities of error. However, this does not lead us to conclude that nobody knows anything and that we should revise our knowledge attributing practices. It therefore seems mistaken to impute to us such an error. Furthermore, this explanation not only commits competent speakers to widespread error but also to systematic irrationality.

---

Footnote 23 continued

that one of the alternative views—either infallibilism or the no-coherent-concept view—offers a better overall theory (i.e. has more explanatory power and few counterintuitive consequences); otherwise the fallibilist might concede that his view has problems and yet maintain that his view is sufficiently better than the alternatives.

Shortly after recognizing that we are fallible, we go right back to saying that we *do* know various things.

Why, then, are we prohibited from saying “I know it is so, but I may be wrong”? If fallibilism is true, why does it seem inappropriate to make ‘concessive knowledge attributions’—i.e., sentences of the form “S knows that  $p$ , but it is possible that  $q$ ” (where  $q$  entails not- $p$ )? To speak of fallible knowledge, of knowledge despite the possibility of error, sounds contradictory.<sup>24</sup> The challenge is to explain why speaking of fallible knowledge is odd without impugning the truth of fallibilism.

Austin again sheds light on the situation. He says that we ought not say that we know if we are aware that we may be mistaken. However,

being aware that you may be mistaken doesn’t mean merely being aware that you are a fallible human being: it means that you have some concrete reason to suppose that you may be mistaken in this case [...] It is naturally *always* possible (“humanly” possible) that I may be mistaken [...], but that by itself is no bar against using the expression “I know”. (1946: 98)

We will credit a person with knowing that  $p$  when she has done what counts as enough in a normal case to demonstrate that  $p$ . Doing enough does not require doing everything. In everyday speech, acknowledging the possibility that one might be wrong occurs when there is some relevant ground for supposing that one is mistaken, and that one therefore isn’t confident that  $p$ . But if that is so, then one shouldn’t claim to know that  $p$ . Concessive knowledge attributions serve to pragmatically impart that the speaker does *not* know the relevant proposition because she cannot rule out all the *relevant* possibilities (i.e. those that the members of the context of utterance ought to take seriously). For example, “It is possible that  $q$ ” conveys that one does not know that not- $q$ . Thus, we should expect “I know that  $p$ , but I might be mistaken” to sound odd.<sup>25</sup> It is obvious for nearly any  $p$  that there is always a possibility that not- $p$ , so Grice’s Maxim of Relation recommends that one not mention it. In virtue of how phrases acknowledging our fallibility usually function, the utterance “I know that  $p$ , but it is possible that  $q$ ” should leave us perplexed (Rysiew 2001, p. 293).<sup>26</sup>

BonJour provides a second explanation for why we typically (but mistakenly) ascribe knowledge in everyday life. He suggests that people may exaggerate when they ascribe knowledge to themselves and to others. Following Butchvarov (1970), BonJour includes the concept of knowledge in a class of concepts that are often used in a way that involves exaggeration. This class of concepts pertains to evaluative stan-

<sup>24</sup> “Possible” denotes epistemic possibility.

<sup>25</sup> Stanley (2005) argues that concessive knowledge attributions sound odd because they are semantically defective rather than pragmatically inappropriate. He writes, “My problem [with the pragmatic explanation] is that it is mysterious to me what he [Rysiew] takes to be the *semantic content* of epistemic possibility statements” (2005, p. 127). However, Dougherty and Rysiew (2009) suggest that what is epistemically possible for a subject are those things which the subject’s evidence does not rule out:  $q$  is epistemically possible for S iff not- $q$  isn’t entailed by S’s evidence. Thus, “I know that  $p$ , but it is possible that  $q$ ” means “I know that  $p$ , but there is a non-negligible probability on my total evidence that  $q$ ”.

<sup>26</sup> Likewise, “Henry is a zebra, but it might be that Henry is just a cleverly painted mule” could express a true proposition and yet sounds contradictory.

dards or ideals that have “serious practical significance” (BonJour 2010, p. 72), an example of which is romantic love. According to BonJour,

It is obvious that romantic love represents a kind of ideally perfect situation with respect to a certain sort of interpersonal relation and also that the applicability of this concept has a range of practical implications that are usually viewed as highly desirable. This creates a strong pressure to attribute the state of being in love to oneself and to others even when the strong (though rather elusive) requirements for being in such a state are not clearly met — and even when they are clearly not met. (*ibid*)

This explanation is also implausible. An infallible concept of knowledge is not more closely tied to evaluative standards that have “serious practical significance”. In everyday life, people do not demand or require that their informants be in a strong enough epistemic position to rule out radical skeptical scenarios. Indeed, it seems clear that no practical purpose can be served by imposing a test that every informant (*qua* fallible human) is bound to fail. It makes no practical difference to our inquiries whether or not a radical skeptical scenario actually obtains because, as far as we can tell, there is no perceivable difference between these two scenarios. If the concept of knowledge is connected to epistemic standards that have serious practical significance, then we should expect downward pressure on the standards required for knowledge.

My view illustrates that we can account for the special status of knowledge as an epistemic achievement in a way that does not imply infallibilism, which vindicates our knowledge-attributing practices. This gives my position a strong advantage over infallibilism.

## 8 The Gettier problem

What about the alleged payoffs of infallibilism vis-à-vis the Gettier problem and the lottery paradox? Having answered BonJour’s two main challenges to fallibilism and having shown that infallibilism is implausible for several reasons, I think we should reject this view regardless of how well it handles these two philosophical problems. Nevertheless, I will also demonstrate that the perceived benefits of infallibilism with respect to the Gettier and lottery problems are in fact illusory.

Let’s start with the Gettier problem. Many philosophers have been persuaded that an agent in a Gettier scenario does not have knowledge, and consequently much ink has been spilt over how to accommodate these cases. By demanding conclusive justification, infallibilism allegedly secures the correct intuitive verdict in Gettier cases. If the luckless fall guy in a Gettier scenario does not know the relevant proposition, then rejecting fallibilism is one way to make this problem disappear. He does not know because he lacks sufficient justification.

However, infallibilism might not escape the Gettier problem entirely. As Lycan (2006, p. 152) points out, this problem resurfaces for the infallibilist in a different shape. Any infallibilist will presumably admit that knowing is at least a regulative ideal, so some cognitive conditions will come closer than others to approximating this ideal. BonJour (2010, p. 73) himself accepts this view; he says that we (falsely)

ascribe knowledge because “a case may seem so close to an ideal standard as to make it seem unreasonably or needlessly fussy to insist on the difference”.<sup>27</sup> If knowing is a regulative idea, then a person who has overwhelmingly strong evidence should be counted as just-about-knowing or knowing-for-all-practical-purposes, even if nobody every strictly knows. However, this holds only so long as the subject is not Gettiered. A Gettier victim “does not just-about-know or as-good-as know; a Gettier victim *simply does not know*” (Lycan 2006, pp. 152–153). This difference remains to be explained by the infallibilist.

Even if we set aside this objection, the infallibilist “solution” to the Gettier problem comes at a high price. To vindicate our intuitions about these problematic cases, the infallibilist imputes to us widespread error in the vast majority of cases in which the word ‘know’ is used in epistemic evaluation. Infallibilism therefore gets the wrong intuitive verdicts across a much wider range of cases. In contrast, fallibilism gets the correct intuitive verdict in the vast majority of cases. We happily and appropriately credit knowledge to people in ordinary life without demanding that they are infallible.

My goal in this section is to play down the importance of the Gettier problem. I take up what Lycan calls the ‘Gettier Problem problem’, which is the task of “explaining what is distinctively wrong with the Gettier project”; i.e. the project of trying to solve the Gettier problem. I do not go so far as to claim that the Gettier problem is pointless, unfruitful, or idle. I think that exploring such cases can help to reveal important insights about our concept of knowledge and its role in human life and thought. However, I will in large part attack the intuitions that generate the Gettier problem. By drawing on empirical research, I will argue that many people do not have the intuitive judgments that generate the problem in the first place. I then try to explain why people who are presumably able to smoothly and effectively employ the word ‘know’ in everyday life nonetheless offer mixed judgments about Gettier cases.

Several recent investigations by experimental philosophers have examined intuitions that people have about Gettier cases with varying results. For example, Weinberg et al. (2001) presented college students from a variety of cultural backgrounds with the following Gettier case:

Bob has a friend, Jill, who has driven a Buick for many years. Bob therefore thinks that Jill drives an American car. He is not aware, however, that her Buick has recently been stolen, and he is also not aware that Jill has replaced it with a Pontiac, which is a different kind of American car. Does Bob really know that Jill drives an American car, or does he only believe it?

Only 26 % of Westerners attributed knowledge in this Gettier case, but a substantial number of East Asians and South Asians (57 and 61 %, respectively) attributed knowledge to Bob. Cullen (2010) repeated this study but asked simply whether Bob knows or does not know that Jill owns an American car, and found that 42 % of his North American participants attributed knowledge to Bob. In Buckwalter’s (2012) study, participants read a Gettier case about a CEO signing some documents and they were asked to attribute knowledge on a Likert scale. Buckwalter found that most partici-

<sup>27</sup> Unger (1975) also defends this view.

pants scored above the midpoint, which suggests that they are inclined to regard the Gettier victim as a knower.

One of the most recent and thorough empirical investigations of Gettier scenarios shows that many people are willing to attribute knowledge in such cases. In ‘The Folk Conception of Knowledge’, Starmans and Friedman (2012) presented people with the following Gettier case (among others):

Peter is in his locked apartment reading, and is about to have a shower. He puts his book down on the coffee table, and takes off his black plastic watch and leaves it on the coffee table. Then he goes into the bathroom. As Peter’s shower begins, a burglar silently breaks into the apartment. The burglar takes Peter’s black plastic watch, replaces it with an identical black plastic watch, and then leaves. Peter is still in the shower, and did not hear anything. (*ibid.*: 274)

Participants were asked whether Peter ‘really knows’ or ‘only thinks’ that there is a watch on the table. Over 70% of people attributed knowledge to Peter in this Gettier case. This finding was confirmed by a similar experiment involving a different scenario (*ibid.*, p. 276).<sup>28</sup>

Moreover, a new study by Nagel et al. (forthcoming) reveals that only 41.1% of respondents immediately denied knowledge in a Gettier case. The scenario they tested involves a jewelry shopper named Emma who selects a diamond necklace from a tray marked ‘Diamond Earrings and Pendants’. Unbeknownst to Emma, however, this store has a very dishonest employee who has been stealing real diamonds and replacing them with cubic zirconium fakes. Emma cannot tell a real diamond from a fake just by looking and touching; moreover, the (real) pendant she selects was one of a very few pendants that did not have cubic zirconium stones. Only 108 of 263 respondents initially denied knowledge to Emma, whereas 155 people (59%) answered “Yes, she knows”. Of the group that initially attributed knowledge to Emma, 23.6% of them were subsequently led to deny knowledge when asked: “In your view, which of the following sentences better describes Emma’s situation: (a) Emma knows that the stone is a diamond; (b) Emma feels like she knows that the stone is a diamond, but she doesn’t actually know that it is”. The remaining 35% answered (a) to the follow-up question and thus had an unwavering knowledge ascription.<sup>29</sup>

These empirical findings are something of a mixed bag, but they do contradict the longstanding assumption that “everybody agrees” Gettier cases obviously do not involve knowledge of the relevant proposition. What should we make of this fact? A significant upshot is that epistemologists are not required to explain why Gettier victims lack knowledge. If Gettier scenarios elicit mixed judgments about whether a subject knows, then an adequate theory of knowledge need not (and should not)

<sup>28</sup> The exceptions were cases in which an agent’s belief was based on apparent evidence—i.e. evidence that only *appears* to be informative about the world. For example, Starmans and Friedman (2012, p. 278) mention a scenario where a student comes to believe that his professor is in her office because the student sees a convincing hologram sitting at the professor’s desk. Participants attributed knowledge more in the Authentic Evidence condition (67%) than in the Apparent Evidence condition (30%).

<sup>29</sup> Nagel et al. (forthcoming) interpret their data to show that Gettier cases are “widely judged to be instances of justified true belief without knowledge”. I do not see how this conclusion is supported by the aforementioned data.

predict that Gettier victims always lack knowledge. Rather, the burden of a fully adequate theory of knowledge would be to provide a plausible explanation for why there is a lack of widespread agreement about these cases.

Conveniently, the method of practical explication can provide such an explanation. We begin by identifying a deep prejudice in philosophy that is encouraged by Gettier scenarios. The prejudice is to concentrate on a particular situation that epistemologists are notably fond of, which Bernard Williams calls the ‘examiner situation’:

the situation in which I know that  $p$  is true, this other man has asserted that  $p$  is true, and I ask the question of whether this other man really knows it, or merely believes it. (Williams 1973: 146)

To illustrate, consider Gettier’s own example in which we are told that Smith has strong evidence for his true belief that the man who will get the job has ten coins in his pocket (call this proposition ‘ $p$ ’). The situation is described such that we know that  $p$  is true, that Smith has asserted that  $p$  is true, and we are to decide whether Smith really knows that  $p$ . In cases of this sort, we are represented as checking on someone else’s credentials for something about which we already know. But this is not our standard situation with respect to knowledge. The central focus of epistemic evaluation is the activity of inquiry. Our inquiries aim at truth, so it is apt that our knowledge ascriptions gauge the progress of inquiry. The examiner situation is concerned with whether some potential knower really qualifies as such, but in the actual business of inquiry one who needs to find a good informant does not know whether  $p$  but wants to. As Williams remarks,

our standard situation with regard to knowledge (in relation to other persons) is rather that of trying to find somebody who knows what we don’t know; that is, to find somebody who is a source of reliable information about something. (*ibid*)

As I have been arguing, there is an important connection between knowledge and identifying people who are good sources of information.

Will solving the Gettier problem do anything to advance or clarify the proper conduct of inquiry? Kaplan (1985) provides a compelling argument to think that solving the Gettier problem—in so far as it *is* still a problem in light of the aforementioned empirical evidence—will do nothing to advance or clarify the proper conduct of inquiry.<sup>30</sup> If you are a responsible inquirer engaged in the process of inquiry, you will carefully consider and evaluate the evidence and then conclude whether or not the weight of your evidence supports  $p$ . Now imagine that you are determining whether someone else knows whether  $p$ . From your perspective, the situation looks different from that of the Gettier victim: you are in a position to tell whether the agent’s belief is justified and yet false, or justified, true, and yet arrived at via a false premise, or else

<sup>30</sup> I do not endorse Kaplan’s conclusion that the concept of knowledge does not provide a useful goal for our inquiries. This conclusion leaves mysterious why we should have cared about knowledge in the first place.

Gettiered in some other way. You can distinguish between various states of affairs that the Gettier victim cannot.

How important is our ability to make these discriminations? This discriminatory capacity does not seem important or relevant to the proper conduct of inquiry. If I can tell that your belief that  $p$  is perfectly justified and arrived at by the best methods of inquiry, but unluckily based on a false belief, then I must already know that  $p$ . Why, then, does it matter to me whether or not the word ‘know’ is appropriately ascribed to you in this case? What difference does answering this question make to my interest in figuring out whether or not  $p$ ? It is not as if I am assessing your own performance as an investigator. It has already been granted that your evidence was admirably gathered and that you did what any reasonable investigator would have done. Your performance was exemplary and thus the propriety of your conduct as an inquirer is not at issue (Kaplan 1985, p. 357). This makes it difficult to see what possibly hangs on solving the Gettier problem from the point of view of inquiry. For the purpose of understanding and advancing rational inquiry, we have no reason to worry about Gettier cases.

Does my account leave mysterious why many people do not regard Gettier victims (or at least some subset of Gettier victims) as knowers? Gettier victims gather their evidence in admirable ways and do what any reasonable, reliable investigator would have done to arrive at a true belief. Thus, Gettier victims seem to be reliable informants in the objectivized sense. So why is it that many people do not regard Gettier victims as knowers? If knowers are objectivized reliable informants, then Gettier victims should qualify as knowers.

I maintain that we have a good reason to regard Gettier victims as unreliable informants, given what we know. Admittedly, the luckless fall guy in a Gettier case is not stripped of credit nor do we cast doubt on his trustworthiness in general. The correlation between having excellent reasons is fallible, but usually unailing. However, the discovery that the correlation between one’s justification and being right about  $p$  was in this instance accidental will effect our attitude toward the informant *in this instance*, quite apart from how it affects our willingness to rely on him for similar issues on future occasions. As Craig notes, relying on such an informant would

produce the retrospective feeling of having run a risk, of having done something that one would not have done had one been just a little better informed at the time, rather like finding that the person who has just driven you 50 miles down a busy motorway without incident hasn’t passed the driving test. (1990: 49)

A Gettier victim can luckily arrive at a justified true belief in a number of ways. When Smith infers that the man who will be hired has 10 coins in his pocket, his reasoning proceeds via a false premise (i.e. that Jones is the man who will be hired). Moreover, he uses a reliable method to acquire a true belief, but that method played no part, or no appropriate part, in his success. Although he used a reliable method, he was not led to a true belief by the features of the method that make it generally reliable. However, it is in the interest of our truth-seeking inquirer to want true beliefs that are not accidental relative to the method used by our informant. It certainly cannot be recommended as a policy that we should rely on reasoning that proceeds via a false premise, as in



the aforementioned cases. In such circumstances it will be lucky that one's reasoning culminates in a true conclusion.<sup>31</sup>

Not all Gettier victims arrive at the truth via a false premise. For example, consider the victim of Goldman's fake barn scenario:

*Fake Barns.* Henry and his young son are driving along an unknown country road. Henry periodically identifies objects in his visual range for the sake of his young son's edification. Unbeknownst to them, however, they are travelling through an area in which all barns but one are papier-mâché façades. Nevertheless, it just so happens that Henry points to the one real barn and says, "That's a barn". (Goldman 1976: 772)

Henry does not obviously reason via a false premise, yet many people believe that he does not know that the object is a barn.<sup>32</sup> However, the truth of Henry's belief again strikes us as lucky in this instance, so it will affect our attitude toward him *qua* good informant in this instance, even if the method he used is generally reliable. The Gettier victim's reliability is understood counterfactually. Gettier victims seem unreliable because they could have easily been wrong had things been slightly different (e.g., Jones might not have had 10 coins in his pocket; Henry might have easily identified a fake barn). There are a number of ways in which a Gettier victim might luckily arrive at a justified true belief. In any such case, however, the informant will seem reliable in one sense (i.e. he uses a reliable method), but unreliable in another sense (i.e. he could have easily been wrong; he was not led to a true belief by the features of the method that make it generally reliable; etc.). This explains why people's judgments about whether a Gettier victim knows are mixed.

These facts explain why people who are presumably able to smoothly and effectively employ the word 'know' and its cognates in everyday life, and who by any reasonable criteria understand and mean these expressions in the same way, nonetheless offer mixed judgments about Gettier cases. There is a discontinuity between Gettier situations and the practices or situations that our everyday use of 'know' serve to regulate. Suppose we encounter a Gettier case in the course of our daily experience. What practical purposes could be served by determining whether Smith (really) knows that the person who has ten coins will get the job (or only believes it)? Why should we need or care to know, or tell, that Smith knows any such thing? The anomalousness of this question should raise doubts about the significance of our answers to it.

<sup>31</sup> For this reason, those writers (e.g. Lehrer 1965; Harman 1973) who sought to solve the Gettier problem by adding a 'no false lemma' condition to the traditional justified true belief analysis certainly had a substantial point in their favor.

<sup>32</sup> Again, there is much disagreement about what "our" intuitions are. The intuition that Henry does not know that he sees a barn has been denied by several philosophers, including Lycan (2006, p. 158), Millikan (1984), Turri (2011, p. 8), and (now) myself. Gendler and Hawthorne (2005) argue that the intuition is unstable, and DeRose (2009, p. 49) says that this judgment is not "clearly enough correct for it to be usable as the premiss of a good argument". There has been no empirical work to support the claim that the majority of philosophers have the intuition that Henry lacks knowledge. Some empirical data shows that non-philosophers do *not* share this intuition (Colaço et al. 2012).

## 9 The lottery puzzle

What about the lottery puzzle? According to BonJour, the lottery example constitutes a simple and decisive *reductio* against the fallible conception of knowledge. In the case of a lottery, we judge that nobody knows on the basis of probability that they will lose, despite the fact that we have incredibly good justification to believe that the agent will lose (i.e. we are aware that the odds of winning are one-in-a-million). Thus, less than conclusive justification seems insufficient for knowledge in the lottery case. This allegedly impugns a wide range of ordinary, commonsense knowledge because we often take people to know on the basis of non-conclusive justification. But if less than conclusive justification were enough for knowledge, we would presumably have knowledge in lottery cases; yet it is intuitively clear that we do not. Fallibilism therefore issues the wrong verdict in this case: no matter how probable it is that somebody will lose, nothing short of conclusive justification will do. If the level of justification required for knowledge need not be conclusive, why do we judge that people lack knowledge in lottery cases?

The fallibilist may handle the lottery puzzle in one of two ways. One approach is to argue that we mistakenly deny knowledge in lottery cases (and hopefully offer some explanation for why this is so). Reed (2010) believes that denying knowledge in the lottery case will lead to skepticism, so he argues that we do know, for each losing ticket, that it will lose. This approach is troublesome for several reasons. If I know that my ticket will lose, why buy it? And if I know that I will lose, why not sell my ticket for a penny? Furthermore, wouldn't parity of reasoning allow me to know that every other ticket will lose?<sup>33</sup> Finally, why bother checking the newspaper for the result of the draw? The claim that I know that my ticket will lose has undesirable consequences. We simply do not know that we will lose.

A better approach is to maintain that we do lack knowledge in lottery cases and try to show why this does not impugn a wide range of ordinary, commonsense knowledge attributions. This approach is more attractive because it does not commit us to the counterintuitive consequences mentioned in the previous paragraph, nor does it commit us to infallibilism (which, as I have argued, is implausible). The challenge is to somehow distinguish between cases in which the object of knowledge is a proposition like "My lottery ticket will lose" and cases in which the object of knowledge is a proposition like "I will be in Cambridge this summer" (which, as mentioned in Sect. 2, seems to entail that I know that I will not die before then). Is there any principled way to distinguish these types of cases?

Here I will avail myself to a plausible solution that has already been proposed by Cohen (1988). I have little to add to Cohen's solution, but I will connect this view with some of my earlier remarks about fallibilism. Whether or not the reader finds Cohen's

<sup>33</sup> Reed supplements this view with an explanation that prevents these different pieces of knowledge from being aggregated into a problematic conjunction (i.e. knowing that tickets 1 and 2 will lose, knowing that tickets 1, 2, and 3 will lose... knowing that the first 999 tickets will lose, so the winner must be ticket 1,000). Reed's (2010, p. 234) solution is to restrict the closure principle to a small number of plausible premises, thereby preventing us from bringing multiple premises into inferential contact. Thus, we can know, for any collection of, say, five tickets, that all of them will lose, but we cannot know that the first 500 tickets will lose. Where to draw the line is difficult, as Reed admits.

solution to the lottery puzzle compelling, I maintain that my version of fallibilism is more plausible than infallibilism for all of the reasons I have discussed in Sects. 3–8.

Cohen argues that we lack knowledge in the lottery case as a result of the statistical nature of our reasons. The statistical nature of the reasons involved in typical lottery cases make the chance of error salient, and the salience of error sets the epistemic standard to the level of conclusive justification. Whether or not the salience of error shifts the contextual parameter for warranted assertability, on the one hand, or for truthfully saying that one knows, on the other hand, will depend on whether one endorses epistemic contextualism (e.g., Cohen 1988) or insensitive invariantism (e.g., Rysiew 2001). I want to continue remaining neutral on this debate. Whether we decide to cast our lot with contextualism or invariantism, the proposed solution is still fallibilist: we can know that  $p$  even though there is a chance of error. When the chance of error is salient, we are reluctant to attribute knowledge; when the chance of error is not salient, we may appropriately (truthfully) attribute knowledge. The chance of error is not salient in ordinary cases where we appropriately say that we know that somebody has lost a lottery because our reasons consist of testimony, a newspaper, or something else that is not explicitly statistical in nature.

By appealing to the mechanisms of context-sensitivity, we can explain away the intuitions that generate the lottery paradox in a way that does not lead us to violate the principle that knowledge is closed under known logical implication. The appropriateness (truth) of a knowledge ascription is sensitive to certain facts that comprise the context. For each context of ascription, there is a standard for how strong one's epistemic position with respect to a proposition  $p$  must be in order to know that  $p$  (or warrantably assert that one knows). The explanation for why the standards are different in lottery cases is that the statistical nature of our reasons makes the chance of error salient in lottery cases, which prevents us from knowing (or appropriately claiming knowledge).

Does this difference really exist? Ordinarily I may come to know that I have lost a lottery by hearing an official television announcement or reading the results in a newspaper. Surely, however, facts about testimony or newspaper reports constitute reasons (evidence) only in conjunction with facts about the reliability of these sources, which is less than perfect. Aren't the reasons more accurately described as, say, the Daily Star reports that  $q$  and that this newspaper is  $n/m$  reliable (where  $n < m$ )? If the reasons are statistical in this way, then the chance of error should be salient on Cohen's explanation. Consequently, we should be reluctant to attribute knowledge.

Cohen's (1988, p. 107) answer is that "we do not normally think of reasons in this way". Our attention is not focused on the explicit statistical probability that the newspaper is wrong. Insofar as we think about the newspaper case, we might imagine a scenario where the information about the winning ticket is transmitted to the newspaper in a typical way: the reporter witnesses the drawing and tells the Editor, who then prints it. Of course newspapers are not perfectly reliable and there is always the chance of a misprint. The point is that the chance of error is not salient in the imagined scenario because our reasons are not described in terms of the newspapers high, but less than perfect, reliability. If we did consider the various ways in which newspapers make mistakes (e.g., misprints), we *would* begin to wonder whether we really did know that

our ticket had lost. But we do not normally think of the case in this way (*ibid*, pp. 107–108).

Why are the standards set in this way? Why does the salience of error prevent us from knowing (or warrantedly asserting that we know) in the lottery case? According to BonJour (2010, p. 68), the “deepest and most serious objection” to any proposed solution to the lottery paradox is the lack of any plausible intuitive rationale for why knowledge would have the proposed requirement. However, it is a general fact about humans that when error possibilities are salient to us, we will deny, or at least feel inclined to deny, that we know. This connects with my earlier discussion about why it sounds inappropriate to say, “I know it is so, but perhaps I am wrong”. Although we are fallible, we tend to acknowledge our fallibility only when we have some concrete reason in the context to think that we might be mistaken. One way error possibilities can become salient is when my reasons are explicitly probabilistic. This does not prevent us from acknowledging in the abstract that our knowledge ascriptions are fallible and yet sometimes true. As Austin (1946, p. 98) says, “being aware that you may be mistaken doesn’t mean merely being aware that you are a fallible human being: it means that you have some concrete reason to suppose that you may be mistaken in this case”. Thus, the explanatory power of this solution is not limited to lottery cases.<sup>34</sup>

## 10 Conclusion

Fallibilists have devoted too little attention to determining what the so-called “magic” level of justification is and why it is important. Anyone who wishes to defend the epistemological doctrine of fallibilism should reflect on what non-conclusive level of justification is needed for knowledge and why this level is supposed to matter. Furthermore, many fallibilists undoubtedly dismiss the Cartesian conception of knowledge far too quickly on account of its skeptical consequences. The fact that we happily bandy about the word ‘know’ in common parlance does not guarantee that any theory falsifying these attributions must be mistaken. Our all-things-considered judgment might recommend that we abandon fallibilism if an infallibilist theory of knowledge offered the best overall theory.

However, this paper has shown that there *is* a plausible way to specify, to a reasonable measure of approximation, the non-conclusive level of epistemic justification that is required for knowledge. My practical explication provides a value-based explanation of knowledge that is not narrowly epistemic, for I locate the value of knowledge in broadly practical, coordinative concerns. Furthermore, I have provided a plausible account of why this level has the significance that is bound up with our understanding of knowledge. I have also provided a plausible intuitive rationale for the further requirements needed to solve the Gettier and lottery problems in a non-ad hoc way. Finally, my view is more faithful to our ordinary practices—to what we say and do, and think it right to say and do, in everyday life.

<sup>34</sup> My solution to the Gettier problem is not available to the lottery puzzle for two reasons: first, I assume there is widespread agreement that lottery ticket holders do not have knowledge merely on the basis of probability (unlike the lack of agreement in Gettier cases); second, lottery ticket holders do not strike us as unreliable in any sense (unlike Gettier victims).

**Acknowledgments** For helpful advice I am grateful to Hallvard Lillehammer, Chris Cowie, Robin McKenna, Nick Hughes, and the participants of a 2013 Cambridge philosophy seminar. This paper was written while I was generously funded by the Social Sciences and Humanities Research Council of Canada.

## References

- Austin, J. L. (1946). Other minds. *Proceedings of the Aristotelian Society*, 148, 148–187.
- Austin, J. L. (1961). *Philosophical Papers*. Oxford: Clarendon Press.
- BonJour, L. (2010). The myth of knowledge. *Philosophical Perspectives*, 24(1), 57–83.
- Brown, J. (2008). Knowledge and practical reason. *Philosophy Compass*, 3(6), 1135–1152.
- Buckwalter, W. (2012). Non-traditional factors in judgments about knowledge. *Philosophy Compass*, 7(4), 278–289.
- Butchvarov, P. (1970). *The concept of knowledge*. Evanston, IL: Northwestern University Press.
- Chisholm, R. (1977). *Theory of knowledge*. Ann Arbor, MI: Prentice Hall.
- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, 2, 91–123.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13, 57–89.
- Colaço, D., Buckwalter, W., & Stich, S. (2012). Epistemic intuitions in fake-barn thought experiments. [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1973351](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1973351).
- Conee, E., & Feldman, R. (2004). *Evidentialism*. Oxford: Oxford University Press.
- Craig, E. (1990). *Knowledge and the state of nature*. Oxford: Oxford University Press.
- Cullen, S. (2010). Survey-driven romanticism. *Review of Philosophy and Psychology*, 1(2), 275–296.
- DeRose, K. (2009). *The case for contextualism*. Oxford: Oxford University Press.
- Dougherty, T., & Rysiew, P. (2009). Fallibilism, epistemic possibility, and concessive knowledge attributions. *Philosophy and Phenomenological Research*, 78(1), 123–132.
- Gendler, T., & Hawthorne, J. (2005). The real guide to fake barns: A catalogue of gifts for your epistemic enemies. *Philosophical Studies*, 124(3), 331–352.
- Gentzler, J. (1995). How to discriminate between experts and frauds. *History of Philosophy Quarterly*, 12(3), 227–246.
- Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23(6), 121–123.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73, 771–791.
- Greco, J. (2000). *Putting skeptics in their place: the nature of skeptical arguments and their role in philosophical inquiry*. Cambridge: Cambridge University Press.
- Greco, J. (2010). *Achieving knowledge*. Cambridge, MA: Cambridge University Press.
- Haddock, A., Pritchard, D., & Millar, A. (Eds.). (2010). *Epistemic value*. Oxford: Oxford University Press.
- Harman, G. (1973). *Thought*. Princeton, NJ: Princeton University Press.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Henderson, D. (2009). Motivated contextualism. *Philosophical Studies*, 142(1), 119–131.
- Henderson, D. (2011). Gate-keeping contextualism. *Episteme*, 8(1), 83–98.
- Hetherington, S. (2001). *Good knowledge, bad knowledge: Two dogmas of epistemology*. Oxford: Oxford University Press.
- Kaplan, M. (1985). It's not what you know that counts. *Journal of Philosophy*, 82(7), 350–363.
- Kaplan, M. (2008). Austin's way with skepticism. In J. Greco (Ed.), *The Oxford handbook of skepticism*. Oxford: Oxford University Press.
- Kappel, K. (2010). On saying that someone knows: Themes from Craig. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Social epistemology* (pp. 69–88). Oxford: Oxford University Press.
- Kelp, C. (2011). What's the point of "Knowledge" anyway? *Episteme*, 8, 53–66.
- Kusch, M. (2011). Knowledge and certainties in the epistemic state of nature. *Episteme*, 8, 6–23.
- Kvanvig, J. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge: Cambridge University Press.
- Kvanvig, J. (2009). The value of understanding. In D. Pritchard, A. Haddock, & A. Miller (Eds.), *Epistemic value* (pp. 95–112). Oxford: Oxford University Press.
- LaBarge, S. (1997). Socrates and the recognition of experts. *Apeiron*, 30(4), 51–62.
- Lehrer, K. (1965). Knowledge, truth and evidence. *Analysis*, 25, 168–175.
- Lycan, W. (2006). On the Gettier problem. In S. Hetherington (Ed.), *Epistemology futures*. Oxford: Oxford University Press.

- Millar, A. (2010). Knowledge and recognition. In D. Pritchard, A. Millar, & A. Haddock (Eds.), *The nature and value of knowledge: Three investigations*. Oxford: Oxford University Press.
- Millikan, R. (1984). *Thought, language, and other biological categories*. Cambridge, MA: MIT Press.
- Nagel, J., Juan, S., & Mar, R. (forthcoming). Lay denial of knowledge for justified true beliefs. *Cognition*.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.
- Pritchard, D. (2009). Knowledge, understanding and epistemic value. In A. O'Hear (Ed.), *Epistemology, Royal Institute of Philosophy supplements* (Vol. 64, pp. 19–43). Cambridge, MA: Cambridge University Press.
- Pritchard, D. (2010). Knowledge and understanding. In D. Pritchard, A. Millar, & A. Haddock (Eds.), *The nature and value of knowledge: Three investigations*. Oxford: Oxford University Press.
- Reed, B. (2010). A defense of stable invariantism. *Noûs*, 44(2), 224–244.
- Russell, B. (1912). *The problems of philosophy*. London: Williams and Norgate.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs*, 35(4), 477–514.
- Rysiew, P. (2012). Epistemic scorekeeping. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions*. Oxford: Oxford University Press.
- Siegel, H. (1997). *Rationality redeemed?*. New York: Routledge.
- Sosa, E. (2003). The place of truth in epistemology. In L. Zagzebski & M. DePaul (Eds.), *Intellectual virtue: Perspectives from ethics and epistemology*. Oxford: Oxford University Press.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, 124(3), 272–283.
- Turri, J. (2011). Manifest failure: The Gettier problem solved. *Philosophers' Imprint*, 11(8), 1–11.
- Unger, P. (1975). *Ignorance: A case for skepticism*. Oxford: Oxford University Press.
- Vogel, J. (1990). Are there counterexamples to the closure principle? In M. D. Roth & G. Ross (Eds.), *Doubling: Contemporary perspectives on skepticism*. Dordrecht: Kluwer.
- Weinberg, J., Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. *Philosophical Topics*, 29(1–2), 429–460.
- Williams, B. (1973). *Problems of the self*. Cambridge, MA: Cambridge University Press.
- Williamson, T. (1994). *Vagueness*. London: Routledge.
- Zagzebski, L. (1994). The inescapability of Gettier problems. *Philosophical Quarterly*, 44(174), 65–73.