

The redundancy problem: From knowledge-infallibilism to knowledge-minimalism

Stephen Hetherington¹

Received: 20 December 2015 / Accepted: 5 April 2016 / Published online: 25 April 2016
© Springer Science+Business Media Dordrecht 2016

Abstract Among the epistemological ideas commonly associated with the Descartes of the *Meditations*, at any rate, is a knowledge-infallibilism. Such an idea was seemingly a vital element in Descartes’s search for truth within that investigative setting: only a true belief gained infallibly (as we would now describe it) could be knowledge, as the *Meditations* conceived of this. Contemporary epistemologists are less likely than Descartes was to advocate our ever seeking knowledge-infallibility, if only because most are doubtful as to its ever being available. Still, they would agree—in a seemingly Cartesian spirit—that *if* infallible knowledge was available then it would be a stronger link to truth than fallible knowledge ever manages to be. But this paper argues that infallible knowledge lacks that supposed advantage over fallible knowledge. Indeed, we will see why we should move even further away from the epistemological model at the heart of the *Meditations*: we should adopt knowledge-*minimalism*, by conceiving of a belief’s being true as always sufficient for its being knowledge—this, for any belief.

Keywords Knowledge · True belief · Knowledge-fallibilism · Knowledge-infallibilism · Knowledge-minimalism · Redundancy problem · Descartes · Sartwell

1 Introduction

There is no more Cartesian an idea within epistemology than knowledge-infallibilism (as a contemporary philosopher might call it). At least within his *Meditations on First*

✉ Stephen Hetherington
s.hetherington@unsw.edu.au

¹ The University of New South Wales, Sydney, Australia

Philosophy, Descartes sought truth by seeking known-by-him truth; and he deemed himself to know a truth only when the associated idea had been gained by him in such a way as to eliminate even the possibility of his mistaking the idea as being true when it is not.¹ In contemporary terms, the Descartes of the *Meditations* was thus a knowledge-infallibilist.

But should he have been so? More strongly, should anyone be so? This paper will argue against one of the more seemingly excellent reasons for conceiving of knowledge in infallibilist terms. Initially, the paper will thereby be advocating a knowledge-fallibilism. Of course, whether there can be fallible knowledge—or whether instead knowledge must somehow incorporate infallibility—depends first on what fallibility is.² What would infallibility give us that fallibility does not, when we are functioning as knowers? In this paper, I discuss one apparent epistemic advantage that infallible knowledge would routinely be taken to have over fallible knowledge. That appearance is badly misleading, I will contend.

2 Distinguishing knowledge-fallibilism from knowledge-infallibilism

Let knowledge-fallibilism and knowledge-infallibilism, at their most generic, be the following theses:³

Knowledge-fallibilism It is possible for at least some knowledge to be fallible.

Knowledge-infallibilism It is impossible for there ever to be fallible knowledge, because knowledge could only ever be infallible.

And what would the difference be between knowledge's being fallible and its being infallible? This difference has traditionally been construed as a function of at least the truth-directedness strength of the justification component within, respectively, any instance of fallible knowledge and any instance of infallible knowledge. Truth-directedness is not the only epistemologically significant possible aspect of epistemic justification, of course. But it has long been the most salient aspect to mention when trying to distinguish just between a fallibilist and an infallibilist sense of justification and thereby knowledge. The central question over which those two approaches differ has typically been one of whether the belief's being *true* would be entailed or, more generally, somehow ensured or guaranteed by the justification, or whether—even

¹ Perhaps Descartes in other guises did not constrain himself in this way. I take the *Meditations* to have been portraying a quest for metaphysical truth in particular. On whether Descartes was content to settle for satisfying a probabilistic standard when seeking truth in general, especially within empirical science, see Clarke (2012).

² On some details that arise for this issue, see Reed (2000, 2002, 2012), Fantl and McGrath (2009), Dougherty (2011), and Hetherington (forthcoming a).

³ Non-modalized versions of these two theses are also available. They would say only that there *is*, or there *is not*, some fallible knowledge. But I take the modalized versions to fit better with the metaphysical, rather than empirical, road to be followed in this paper. (I call that road metaphysical, not conceptual, because the paper's project is not the traditional one of seeking necessary and sufficient conditions for the satisfaction of a concept. The emphasis will instead be on what is involved, essentially versus accidentally, or inherently versus extrinsically, in a given belief's being an instance of knowledge).

though the belief is true—its being false was in some way allowed by the justification (Hetherington 2005, forthcoming a, b).

Upon epistemological analysis, that central question readily becomes more complicated, such as by asking about whether the belief is *accidentally* true, relative to the justification (Reed 2000, 2002, 2012), or about whether the belief is *failably* knowledge (Hetherington 1999, 2001, Chap. 2), for example.⁴ Still, throughout all such variations on knowledge-fallibilism's central theme, the following minimal component recurs—the idea that there is some sort of compatibility (whatever form, more exactly, that compatibility takes) between (i) the justification within the knowledge (whatever form, more exactly, that justification takes) and (ii) the belief's not being true (no matter that it is actually true).

And that recurring minimal component, talking as it does of compatibility and truth, admits of being parsed modally. Doing so produces, I will assume, a characterisation that includes something relevantly like this:

Within at least one possible world from within whatever group of possible worlds is most apt for modelling the fallibility or otherwise of a particular true belief's being justified, the belief is false—no matter that the belief is in fact knowledge.

The generality in that formulation will enable the paper's argument to be adverting, with systematic ambiguity, either to metaphysically possible worlds or, more narrowly, to epistemically possible ones—without our needing to choose in this setting between these respective ideas as ways of understanding the fallibility or otherwise of some knowledge.⁵ That difference will be immaterial to the argument. In *either* case, the point is still this: we have ready to hand a philosophically congenial means of tak-

⁴ Incidentally, these ideas are general enough to accommodate either *internalist* or *externalist* accounts of justification. For a formulation of this epistemologically significant distinction, see Hetherington (1996, Chaps. 14, 15). On an internalist account, the question is one of whether the person's *evidence*—where, contrary to Williamson (2000), I do not presume this to be knowledge—has a content that entails the truth of the belief. On an externalist account, the question is one of whether the obtaining of some further *circumstance*—such as the belief's having been formed in a truth-conditionally reliable way—is compatible with the belief's nonetheless being false.

⁵ In using the term 'epistemically possible', I am not presuming any specific full conception of epistemic possibility. For example, suppose that we use the term as Hintikka (1962) did, to mean 'compatible with what one knows'. A world *W* would thus be epistemically possible, relative to one's knowledge within this world, if and only if whatever obtains in *W* is compatible with whatever one knows within this world. But that conception will not tell us about fallible knowledge's difference from infallible knowledge. It implies that if in this world one knows either fallibly *or* infallibly that *p*, then no epistemically possible world for one is a world where not-*p* obtains. Yet—on the core commitment (mentioned above) behind the traditional idea of fallible knowledge that *p*—*somewhere* there is at least one *somehow* relevantly possible world where, when one's knowledge that *p* is fallible but not when it is infallible, not-*p* obtains. This remains so if we think of epistemic possibility in other ways. Suppose that not-*p* is said to be epistemically possible for one even when one knows that *p*, in case one does not know that not-*p* is not compatible with what one knows. This still allows not-*p* to obtain in at least one world that would help to model such a possibility; for not-*p* is being said to be compatible with one's *bi*-level epistemic stance on *p* at that time. (Hintikka [ibid., Chap. 5] argues that knowing entails knowing that one knows, and so would not allow this possibility.) Or suppose that, like Chalmers (2002), we say that not-*p* is epistemically possible, even when one knows that *p*, if it is not ruled out *a priori*: it is not known *a priori* that *p*. Explicating this, too, will call upon the existence of at least *some* relevant not-*p* world(s). For more on how to conceive of fallible knowledge in terms of possible worlds, see Hetherington (forthcoming a).

ing at least a first step—even if a schematic and programmatic one—towards parsing modally the epistemologically traditional idea of there being a difference of justificatory *strength* between any instance of fallible knowledge that p and any instance of infallible knowledge that p . Specifically, the modal translation of that sort of claim about a difference of justificatory strength will tell us this:

There is at least some salient not- p possibility that any fallible knowledge that p 's justification does—while no infallible knowledge that p 's justification does—*fall short* of eliminating.

In effect, we are being told that, for any given instance of knowledge that p (and so long as all else is equal), *at least one more* group of possibilities—from among all of the relevant not- p ones—is eliminated if the justification within the knowledge that p is infallible than if the justification within the knowledge that p is fallible.

The availability of some such line of thought matters, because it is seemingly an epistemological truism—a claim accepted without hesitation by epistemologists—that infallible knowledge that p would somehow be a stronger justificatory link to the truth that p than fallible knowledge that p would be. Let us parse the core of that standard idea more fully and precisely, still in terms of possible worlds.⁶

KInF > *KF* For any given proposition p , consider both (i) any possible *fallible* justification ${}_pJ_F$ within this world α for a belief that p and (ii) any possible *infallible* justification ${}_pJ_I$ within α for a belief that p . For any such ${}_pJ_F$, let the ${}_pJ_F$ -worlds be those where the belief that p is formed on the same fallible justificatory basis ${}_pJ_F$ (the same justifying evidence and in the same justifying circumstances that jointly constitute ${}_pJ_F$) as does or could occur for that belief within α . For any ${}_pJ_I$, too, let the ${}_pJ_I$ -worlds stand analogously to ${}_pJ_I$: that is, they stand to ${}_pJ_I$ as, for any ${}_pJ_F$, the ${}_pJ_F$ -worlds stand to ${}_pJ_F$. Let ${}_p\%T_F$ be the proportion of accessible p -worlds (those where it is *true* that p) among those ${}_pJ_F$ -worlds; and, analogously, let ${}_p\%T_I$ be the proportion of accessible p -worlds among those ${}_pJ_I$ -worlds. Then ${}_p\%T_I > {}_p\%T_F$ —because ${}_p\%T_I = 100\%$ while ${}_p\%T_F < 100\%$. Next, assume that any fallible *knowledge* that p would include some ${}_pJ_F$ and that any infallible *knowledge* that p would include some ${}_pJ_I$. Then we may infer that no fallible knowledge that p includes within itself as strong a justificatory link to the truth that p as does any infallible knowledge that p —again, because ${}_p\%T_I > {}_p\%T_F$.

⁶ Doing so does not preclude the possibility that the disparity between those respective justificatory strengths would be better articulated without mentioning possible worlds at all. But in this respect the explicative onus, I suggest, is upon those who would seek to evade this approach. The difference in justificatory strength has been explicated above in modal terms, by talking of the presence of a *possibility*—a compatibility—that accompanies fallible but not infallible justification. Accordingly, until I am aware of a better way of modelling possibilities, I will continue speaking of possible worlds in this setting. (In this respect, too, I take heart from Pritchard's (forthcoming) argument that we should conceive in modal terms of the related aspects of knowledge.)

3 The redundancy argument

But is that epistemological truism— $KInF > KF$ —actually true? Here is an argument against its being so.

Let FJTB be the property of being a fallibly justified true belief, and let InFJTB be the property of being an infallibly justified true belief. Included within each of those complex properties is the property T, the property of being true. Hence, in no possible world is there an instance either of FJTB or of InFJTB that is not an instance of T. Obviously there are worlds where neither of those two complex properties—FJTB and InFJTB—is instantiated, either for a given belief that p or even at all. Yet in no possible world where InFJTB is instantiated is the belief in question not true. And, likewise, in no possible world where FJTB is instantiated is the belief in question not true. So (for a given p), $p\%T_I = 100\%$ and $p\%T_F = 100\%$. In short, $p\%T_I = p\%T_F$. And thus thesis $KInF > KF$ is false: the justificatory support for p even within some infallible knowledge that p is *not* a stronger link to its being true that p than is the justificatory support for p within some fallible knowledge that p.

I call that *the redundancy problem*. It identifies a respect in which infallible justification for a true belief is always redundant, once there is fallible justification for that true belief. Initially, we are encountering this as a potential problem for infallibilist conceptions of knowledge.⁷ It tells us that *no* higher a proportion of relevant possibilities of falsity is eliminated by infallible justificatory support for a true belief than is eliminated by fallible justificatory support for that same true belief. In that sense, fallible justificatory support for a true belief *may as well remain* fallible: no proportional improvement in justificatory strength could be achieved by replacing the fallible support for that true belief with infallible support for that same true belief. Apparently, therefore, this is a sense in which infallible justificatory support for a true belief would be no stronger than fallible justificatory support for that true belief. And so—given knowledge's needing to include justificatory support that is either fallible or infallible—we find that even infallible knowledge that p would be no stronger a justificatory link than fallible knowledge that p would be to the truth that p.

4 The swamping problem

Section 3's redundancy problem is a close cousin of what epistemologists call the *swamping* problem.⁸ The latter arose initially as a putative objection to process-reliabilist accounts of knowledge. In that form, the swamping problem talks of epistemic value, and it denies that there is any added epistemic value in having a reliably acquired true belief that p, beyond whatever epistemic value there is in having a true belief that p. The thinking behind the problem is as follows:

⁷ In Sect. 8, however, we will extend this section's argument, by investigating how the redundancy may be generalised so as to pose a challenge also to knowledge-fallibilism.

⁸ For some of the swamping problem's history, see Riggs (2002), Kvanvig (2003), Zagzebski (2003), Olsson (2011), Davis and Jäger (2012), Bates (2013), and Dutant (2013).

The process-reliability of which many epistemologists speak so favourably is a truth-directed reliability: it is a matter of how reliable a given belief-forming process or method is in generating true beliefs, at least within actual situations but perhaps also within alternative possible ones. Now, imagine a person's *having* a true belief via a reliable belief-forming process. How is that true belief's presence epistemically better because the process—due to its reliability—could easily have produced *further* true beliefs, say? It is not. (Nor is the true belief's reliable pedigree to be respected purely for its *own* sake.)

Clearly, there is a similarity between that line of thought and Sect. 3's reasoning for the redundancy problem. Nevertheless, they remain distinct problems.

First, the reasoning for the redundancy problem does not rely upon the thesis that generates the swamping problem and that Olsson (2011, p. 175), following Goldman (1999), calls Veritism: 'All that matters in inquiry is the acquisition of true belief.' Section 3's argument was about $KInF > KF$. What was being discussed there was a comparison that concerned just the supposed *increase in justificatory strength* between, respectively, fallible knowledge and infallible knowledge. That there would be such an increase between those two states will not be considered epistemologically controversial. It is far less controversial, at any rate, than Veritism's claim that the *only possible epistemic value* involved in inquiry is that a true belief is gained. The 'justificatory increase hypothesis' (which I adopt here for the sake of argument, since others adopt it) reflects the epistemologically standard view that—regardless of whatever will not change between knowing fallibly that *p* and knowing infallibly that *p*—there has to be *some* sort of epistemic strengthening of a believer's standing in relation to *p*, whenever she knows infallibly rather than fallibly that *p*. The usual way to explicate this epistemic strengthening does not need to talk of epistemic values as such; and hence Sect. 3's reasoning for the redundancy problem has not done so. All that has entered the reasoning is a modal way of modelling—not of valuing—that difference of epistemic strength, a difference that epistemologists would typically think is at least part of what is involved in the difference between knowing fallibly and knowing infallibly.

Second, as the next section will explain, the argument for the redundancy problem need not be regarded as being about how the true belief has been *formed* or *produced* (such as in a reliable way).

5 The justificatory link to truth

I have presented the redundancy problem as telling us that no stronger a justificatory link to a belief's being true is constituted by having infallible, rather than fallible, justificatory support for the belief's truth. Yet how could such a thesis about justification and truth be correct? For a start, it clashes with the usual view that infallible justificatory support is a perfect justificatory link to the particular truth, while fallible justificatory support is only an imperfect link to that truth.

That usual view, however, might be overlooking the following distinction:

(1) *Active justificatory linking* This is a process of justifying. It is one's *gaining* a true belief, or at least one's *trying* to do so, by the use of what would *then*—if one

was to proceed to gain that belief—constitute one’s justification for the belief’s being true. Here, we are asking whether one’s justification—be it fallible or be it infallible—*will* lead one to a true belief.

(2) *The state of being justificatorily linked* This is one’s *having* a true belief, perhaps—but not necessarily—as a result of a prior (and, in the sense described just now by (1), an) active justificatory linking. Here, we are asking whether one’s justification—be it fallible or be it infallible—*has* led one to a true belief.

Presumably, an active justificatory linking to a truth is what an inquirer seeks, as she moves to and fro between data, observations, stray thoughts, hypotheses, beliefs, etc.—with all of this being intended to lead her to a state in which she is linked justificatorily to a pertinent truth. In contrast, however, the state of being justificatorily linked to a truth is what needs to be at least our initial focus as we analyse epistemologically what knowledge *is*—which is to say, what it is *to be in* a state of knowing. The redundancy problem arises just for the latter—that is, for the state of knowing, and thereby for the state of having a true belief that has been justified either fallibly or infallibly. The redundancy problem’s question is then the following one. *Once one has* some knowledge that *p*—or, equally, *in having* that knowledge—how could infallible justification within that knowledge ever be contributing a stronger justificatory link, beyond what fallible justification within the knowledge would be contributing, to the truth that *p*?

I am not thereby asking whether, given the redundancy problem, fallible justification and infallible justification are equally good in all epistemic respects—even in all justificatory respects. In particular, I acknowledge, they are not equally good *active* justificatory links to truth. If an inquirer gains evidence that infallibly justifies *p*’s being true, whereupon she forms the belief that *p*, then she has been actively linked to a true belief in a way that guaranteed her belief’s being true once formed. An inquirer actively relying instead upon fallible justification for *p*’s being true has no such guarantee, as she contemplates whether to proceed to believe that *p*. Perhaps her active justificatory linking will in fact be to that truth; right now, though, there is no guarantee that this will occur. No matter: we have begun to see why this sort of justificatory disparity is in any case not the sort of situation that is being discussed by Sect. 3’s argument for the redundancy problem.

Let me expand upon that point. A corollary of the redundancy problem is that whatever increased truth-likeness there is in an infallibly justified belief (as compared with a fallibly justified belief) could only have been present at an *earlier* stage of the belief-forming process, not insofar as we are comparing the two relevant *completed* states with each other—the states of being a fallibly justified true belief and being an infallibly justified true belief. Hence, if we insist on there being a stronger link to the truth that *p* in having an infallibly justified true belief that *p*, then—given the redundancy problem—we are pointing only to what would have been the greater likelihood, given some infallibilist justificatory support at an *earlier* moment of the overall justifying process, of a true belief that *p*’s *proceeding to come into* existence. And this increased truth-likeness is part only of an *active* justificatory linking, relative to that particular belief. So, it cannot still be present once the true belief *has* eventuated: after all, at that stage the potential for a true belief to eventuate—again, a potential that

would have been stronger, in advance, within any infallible justification than within any fallible justification for that potential belief—has been *realised*. At which stage, however, the resulting state—the justified true belief—must be assessed against a *different* kind of criterion for strength-of-justificatory-link-to-truth. Specifically, at that stage it has to be assessed as a *static* (rather than an active) justificatory link to the truth that *p*. And the verdict, according to the redundancy problem, is that in this respect the resulting state *may as well* have been produced in an infallible way. Yes, this is an assessment ‘after the event’. But the point is that the justificatory link being assessed at that stage is itself the *state* that has *been* produced, not the *process* that, earlier, might—or might not—*proceed to* produce that state.

6 Unsafety and veritic luck

We can apply Sect. 5’s general point to a case study. Recent epistemology has included much discussion of the concept of epistemic safety, including the idea that a belief is not knowledge if it has been formed in an epistemically unsafe way.⁹ This approach amounts to suggesting (as Mylan Engel and Duncan Pritchard, notably, have done) that even a true belief is not knowledge if it is true only in a *veritically (epistemically) lucky* way;¹⁰ which is to suggest a condition along these lines:

A belief is not knowledge if it is false within too many of the closest possible worlds where the belief is formed in the same way (hence, in particular, on the basis of the same evidence) as it is within this world.¹¹ (For convenience, I will call those worlds *justification-mirroring truth-failure* possible worlds for the belief in question.)

I mention this sort of condition because seemingly it describes a truth-linked justificatory respect in which, many epistemologists would presume, a fallibly justified true belief that *p*—but not also an infallibly justified true belief that *p*—*could* be less strong as a link to *p*’s being true. That comparative weakness would be explicable in modal terms. Most epistemologists will allow that whenever a true belief is justified only fallibly—in contrast to when it is justified infallibly—there was also at least a chance of its being formed unsafely; in which event, it would have been true only in a verit-

⁹ See, for example, Sosa (1999, 2007), for seminal advocacy of this conception.

¹⁰ See Engel (1992, 2011) for the initial such use of the term ‘veritic luck’, and Pritchard (2005, 2007, 2009, 2013) for refinements of the concept. Pritchard (forthcoming) now argues that we should focus epistemologically on veritic (epistemic) *risk* rather than veritic (epistemic) luck. Even so, I will present this section’s discussion in terms of veritic (epistemic) luck, since at this stage it has become such an epistemologically familiar topic. But in any case I believe that the section’s discussion will also apply, *mutatis mutandis*, to the concept of veritic risk.

¹¹ Perhaps more generally, here is an alternative version of that condition:

A belief is not knowledge if it is false within too many of the closest possible worlds where the belief is formed in the same way and within the same circumstances as it is within this world.

The argument I am about to present will not rely upon any specific version of the general idea behind this condition.

ically lucky way.¹² Thus, let W be one of those justification-mirroring truth-failure possible worlds, for a given belief that p :

W is a not- p world, and its existence is (in accord with the generic account of veritic luck mentioned just now) part of the modal dimension of the given belief that p 's being true in *this* world α in a veritically lucky way.

Within W , the fallibly justified belief that p reappears, formed in the same (mirroring) way as within α , where it is likewise fallibly justified. Within W , however, the belief is false—unlike in α . I expect that many epistemologists would then accept the following thesis:

If the belief that p had instead been justified *infallibly* within α , there would be no such W , where it is not true that p . So, at least *that* additional truth-directed justificatory strength—a modally explicable form of strength—is part of a true belief's having been infallibly rather than fallibly justified.

However, that reasoning would be invalid, as I will now explain.

The reasoning begins by positing a justification-mirroring truth-failure world W where the pertinent way of forming a belief (let us call that way *WAY*) has led to a false belief that p —instead of a true one, as has happened in α . We are asked to consider the existence of W as being, in that respect, a modal manifestation of the truth-directedness weakness within the (fallible) justificatory link that is constituted within α by the use of the fallible *WAY*. Now, within α , *WAY* has in fact led to a true belief that p —whereas, within W , the result of using *WAY* is a false belief that p . The proposed (standard) reasoning then takes into account *both* of those worldly points about *WAY*, in deciding that *WAY* has led only unsafely within α to the true belief that p . That is, *WAY* is being considered in a *transworld* manner, even as part of evaluating its justificatory strength just within *this* world. That much is epistemologically familiar fare.

However, that epistemologically familiar line of thought can show only that *WAY*'s modally explicable weakness as a justificatory link to the truth that p is a weakness within it as a fallible *and active* linking to any given true belief that p . This point is less familiar, but it is easily seen (as follows). The transworld explicative structure that is being portrayed reveals only *WAY*'s being a modally less-than-wholly-reliable means of *coming* to have a true belief that p within various worlds: in using *WAY*, one will *come* to form a belief that p ; and, in doing so, one might (as within α)—but also one need not (as within W)—come to have a *true* belief that p . (It is precisely the fact that within W the belief being formed is not true that reveals this evaluation of *WAY* as not succeeding in being an evaluation of *WAY* *after* *WAY* has produced

¹² Sosa has also argued (2011, pp. 84–85) *against* conceiving of knowledge in terms of safety, in favour perhaps of letting aptness—'manifestation of competence'—be the pertinent explicative phenomenon. He also allows that '[a]ptness comes in degrees' (ibid., p. 10) and hence that a belief could manifest a less-than-wholly reliable competence, thereby being formed aptly yet fallibly. This apt formation of the belief could be *fully* apt, too, by manifesting meta-aptness (such as a competence in choosing to exercise the competence that has been manifested aptly in the belief's formation). Nonetheless, I am not sure that this recourse to talk of competence-manifestation *will* accomplish that explanatory goal in a way that allows knowledge still to be fallible. For detailed discussion of this point, see Hetherington (2016, Sect. 5.4).

the true belief within α .) That sort of truth-directedness modal weakness is absent, of course, when the relevant way of proceeding to form a belief is an *infallible* and active link to the truth that p . Yet the existence of that sort of disparity was already conceded in Sect. 5, where we observed its doing nothing to undermine the impact of the redundancy problem. For that problem, we noticed, applies only to justification as a *static* link to truth, not as a means of *coming* to have a belief (which might turn out to be a true belief).

Correlatively, nothing in that standard thinking about WAY and its potential to be used actively both within this world α and within a justification-mirroring truth-failure world W shows that, *once the active justificatory linking by WAY has been completed within α* (with a true belief that p now being in place), *the resulting state*—the fallibly justified true belief that p —is a modally weaker justificatory link to the truth that p than would obtain if an infallible active linking had instead been the cause within α of the true belief that p 's coming to be. After all, the fact remains that in no possible world is a fallibly justified true belief not true—*just as* there is no possible world where an infallibly justified true belief is not true. This fact remains, as a fact about *being in* those respective (completed) states.

Once more, therefore, we meet the redundancy problem. When an advocate (such as Pritchard) of the explanatory efficacy of the concepts of epistemic safety and veritic luck asks us to consider a *not- p* world W where the belief that p has been formed in the same way—WAY—as within α (a p -world), this amounts to asking us to be using W as part of our evaluating WAY's strength as a justificatory link to the truth that p *only insofar as WAY has not yet* resulted (within any given world, such as α or such as W) in a true belief that p . (The hypothesized advocate is asking, in effect, whether using WAY *will* produce a true belief; and we proceed to say that in W it has not done so.) Correlatively, the usual epistemological thinking on behalf of that evaluation of WAY can amount (even if unwittingly) only to highlighting the fact that, *considered in this active respect*, the use of WAY might—but need not—proceed to generate a true belief that p : in α it does so, whereas in W it does not. Along such lines, therefore, the epistemologically standard thinking amounts only to evaluating WAY as a putative *active* justificatory link to truth; which, we saw in Sect. 5, is beside the point that is being made by the redundancy problem. Consequently, that problem survives the proposed (and epistemologically standard form of) counter-argument that this section began by proposing.

7 The Gettier problem

Upon being confronted by the redundancy problem, a further objection that would occur readily to many epistemologists is this:

Whenever a true belief is justified only fallibly, the justificatory door remains open for it to be part of a *Gettier case*—that is, to be *Gettiered* (and thereby not knowledge).¹³ Even if in fact the belief is not Gettiered, the *potential* for

¹³ See Gettier (1963) for the original instances of this phenomenon—the first two *Gettier cases* (as they speedily became known). A belief is *Gettiered* when it is the central belief within a situation saliently similar

its being so was present, since there was fallibility in the justification on the basis of which the belief has been formed. Conversely, whenever a true belief is justified infallibly, that door has been closed, since there are no infallibly justified Gettiered beliefs.¹⁴ This difference may readily be thought of as just one further reflection of the fact that infallible justificatory support is a stronger link to truth than fallible justificatory support manages to be.

Such a suggestion calls upon the widespread epistemological conviction that a belief's being Gettiered precludes its being knowledge. Nevertheless, as I will now explain, what is described by the suggestion is *not* the existence of a stronger justificatory link to a truth that *p* within the state of being an infallibly justified true belief that *p*. This is because the redundancy problem reappears (*mutatis mutandis*, and as follows) for Gettier cases and the suggested way of thinking about them:

Let *G* be the property of being Gettiered, a property categorially applicable to beliefs.¹⁵ *G* is a complex property. It includes at least FJTB—the complex property of being a fallibly justified true belief¹⁶—and hence the property *T* (of being true). Consequently, because *G* includes *T*, there is no possible world where a belief is Gettiered without being true. Yet—via the following reasoning—this also renders the relationship between *G* and *T* relevantly like that between InFJ (the property of being *infallibly* justified) and *T*:

Let *b_p*—a particular belief that *p*—be infallibly justified. Then *b_p*'s being justified in the way it is entails *b_p*'s being true: InFJ includes *T*.

Let *b_p* be Gettiered. Then *b_p*'s being justified in the situation it is in (that is, within the surrounding Gettier case as such) entails *b_p*'s being true: *G* includes *T*.

Epistemologists say, following Gettier himself (1963, p. 121), that the justification within a Gettier case is providing only *fallible* support: the property *G* includes the property *FJ*, just as it includes the properties *T* and *B*. Even so, that standard way of speaking could mislead us into thinking that fallible justification is a *weaker* link than infallible justification is to truth (with Gettier cases

Footnote 13 continued

to the ones described by Gettier. Any such belief is true and justified fallibly; yet it is also, according to most epistemologists, not knowledge. For overviews of the history of epistemological engagement with the challenges for understanding knowledge that were spawned by Gettier, see Shope (1983) and Hetherington (2011). For extended critical discussion of that history, see Hetherington (2016).

¹⁴ For elaboration of this point, see Zagzebski (1994) and Howard-Snyder et al. (2003).

¹⁵ This is the conceptually *primary* range of applicability for *G*. Secondary applications of it are as follows. A *person* is Gettiered insofar as she has a Gettiered belief. A situation is a Gettier *case* insofar as it is centred upon a Gettiered belief (and thereby upon a thereby Gettiered person).

¹⁶ FJTB would routinely be assumed to include more besides, if only as details *within* those individual components *F*, *J*, *T*, and *B*. Many of those possible details attract continuing epistemological debate. (For example, this is where a requirement of the belief's being formed *safely*—the condition described in Sect. 6 above—would be held by some epistemologists to enter the story in an explicative role.) Epistemologists also standardly claim that instantiating *G* precludes being knowledge, so that *G* cannot be even partly coextensive with the property *K*.

being an exemplification of this moral). Yes, the justification for b_p , considered in itself, is instantiating FJ—and does not thereby include T. Even so, G includes T; and hence no belief can instantiate G (thereby instantiating FJ, too) without instantiating T. Consequently, G *may as well* be including InFJ rather than FJ, insofar as the enclosed (static) justificatory link to the truth that p is concerned. Perhaps surprisingly, then, we find that being Gettiered—and thereby being fallibly justified in that notorious way—is *not* a state with a weaker justificatory link to the relevant truth than a state built around infallible justification would be.

The point, more simply, is this. The redundancy problem (as it has been generalised in this section) tells us that, once the state of being Gettiered is present *as a whole*—that is, once a belief *is* Gettiered—the justificatory link within that completed state may as well be reflecting the Gettiered belief's having instead been justified infallibly within that circumstance. The justificatory link between being Gettiered—that is, the state of *being* Gettiered, of *now being in* that complex state as a whole—and being true is as strong as that.¹⁷

8 Knowledge-minimalism

So far, we might regard the redundancy problem as amounting to an argument for a kind of knowledge-*fallibilism*, since it aims to undermine what would, for many, be a potential motivation for seeking to be a knowledge-*infallibilist* in preference to being a knowledge-fallibilist. Even so, we have not yet confronted the redundancy problem's full potential significance. We now need to notice how an instance of it arises also for knowledge-fallibilism.

Thus, compare a fallibly justified true belief that p (this being at least part of what is described in any standard knowledge-fallibilist conceptions of knowledge that p) with what epistemologists call a mere true belief that p. More precisely, compare *instantiating the property FJTB* with *instantiating the property TB*. It is no more possible to instantiate TB without instantiating the property T than it is to instantiate FJTB without instantiating T: in neither case is it possible at all. Accordingly, there is the same sort of redundancy relationship between FJTB and TB as that which was described earlier

¹⁷ Perhaps this can help us to understand an epistemologically *disjunctivist* line of thought that has recently been advocated, in particular by McDowell (2011). For example, he says that one can gain '*indefeasible* warrant for perceptual beliefs' (ibid, p. 38). Suppose that you seem to be confronted by something green. In that event, your 'perceptual state leaves no possibility that [the thing] is not green' (ibid.), *given* that your perceptual state has resulted from 'a non-defective exercise' (ibid, p. 39) of your 'capacity to tell whether things in [your] field of vision are green' (ibid, p. 38). In other words, once there *has been* an exercise of the relevant perceptual capacity, what amounts to a kind of infallibility is present—a kind that fits with the picture I am drawing here. In that respect, it is significant that McDowell regards himself as having provided a *fallibilist* picture; for, he says, although capacities can misfire, his focus is on those times when they do not. Notice how this account is like the way in which, just now, I have described Gettier cases: even while there is fallible justification within them—in the sense of a fallible *active* justificatory link having been used by the person who forms the Gettiered belief—the *resulting* state, of the belief's *now being* in the state of being Gettiered, is itself a watertight static link to the truth in question. (But I will not develop this comparison any further in this paper).

as obtaining between InFJTB and FJTB. And so, if there is a redundancy problem for InFJTB in relation to FJTB, there is one likewise for FJTB in relation to TB. Previously, the redundancy problem told us that any *infallible* justification as such that is present as part of an instantiation of InFJTB is redundant—in the sense of providing no extra strength—as a link to truth. Now, the redundancy problem has become broader in its thinking: now, it delivers the same verdict even for the *fallible* justification within any instantiation of FJTB.¹⁸

Most epistemologists, I expect, will wish to treat such a verdict as a *reductio* of the reasoning behind the redundancy problem. They will say that, if anything is clear within epistemology, it is that any belief that *p* is less strongly linked to the truth that *p* insofar as it is a *true* belief than insofar as it is a *true and justified* belief. After all, this conviction is why we require justification at all within knowledge. Is the redundancy problem's thinking therefore mistaken?

I do not believe so, because I regard that standard epistemological reaction as being needlessly restrictive in its conception of the relationship between knowledge and justification. Obviously we may readily allow that there is a *justificatory* strength in having some justification for one's true belief, and hence within any instance of knowing that includes some justification. But that is hardly surprising: justification is justification, regardless of whatever else it is. Beyond acknowledging that triviality, though, we should wish to ascertain more substantively what further kind of truth-linked strength, if any, there is in the justification's presence. Epistemologists have had much to say about this. Still, we need not engage here with that body of writing, because there is a conceptually prior issue to be confronted—as follows—about the metaphysics of knowing.

I assume that we are standardly being told that, whenever there is fallible knowledge that *p*, the associated justificatory strength (1) would be within, or part of, the knowledge (because the fallible justification is literally a component or part of the fallible knowledge), and (2) would be explicable as a correlatively stronger link to the truth in question (stronger than the link that would be present if the true belief was present without also being justified). But the redundancy problem (as it has been generalised further in this section) may then be interpreted as questioning whether that combination of (1) and (2) is how we *need* to understand the location and role of the justificatory strength associated with an instance of knowledge.

First, here is the relevant application of the generalised redundancy problem:

The presence of the justification—be it fallible or be it infallible—within some knowledge that *p* serves no purpose towards the knowledge's being a stronger link to the truth that *p* than would be served by the presence simply of the true belief that *p*. As regards being a state, or part of a state, that is linked to the particular truth that *p*, the true belief that *p* *may as well* be accompanied by its being unjustified as by its being justified.

¹⁸ Strictly, this verdict is gained by a *limiting* case of the relevant form of thinking, since TB's link to T is not justificatory at all. Nonetheless, modally speaking, it is still the same form of (static) link to T as was the (static) justificatory link from FJTB and from InFJTB. In order to avoid confusion, I will now talk more simply of a (static) link to truth—without still using the term 'justificatory'.

Next, notice that even this need not lead (as would standardly be assumed to follow from the redundancy problem's thinking) to our *dismissing* the importance, for knowing that *p*, of having justification for the belief that *p*. For there is a structurally available escape from that conclusion. Specifically, we may discard thesis (1), as the price to be paid for retaining thesis (2); *and we may come to understand why it is not so high a price*. More fully: we may react to this further application of the redundancy problem by recognising that there *is* still a way for knowledge that *p* to be associated, via justification, with a stronger link to the truth that *p*. Our first step towards this alternative interpretive framework is to *cease requiring that justificatory link to satisfy (1)*. What would this mean, as a resulting view of knowledge? By discarding (1), we would be conceiving of each instance of knowledge in such a way as to allow us to place any associated justificatory link *outside* the knowledge (which would itself be interpreted as at least a true belief). This would allow us to bypass the redundancy problem, which is generated by (1)-plus-(2). And then—this is our second interpretive step here—we will be able to retain (2): we will be able to *return* to accommodating—even if non-standardly (as we are about to see)—the justification's functioning as a stronger link to truth.

How would that be so? Can we interpret that suggested combination—not-(1) plus (2)—less schematically? Indeed so: we need only to call upon Sect. 5's distinction between two ways of being linked to a truth. The present interpretive aim is to accommodate, in a constructive way, this implication of the redundancy problem:

If the justification associated with a case of knowledge that *p* is somehow to be a further and stronger link to the truth that *p* (a stronger link beyond what is effected by the true belief that *p*)—so that the redundancy problem is to be bypassed—then the justification needs to be playing that role *not* from within the knowledge that *p* as such—that is, not as *part* of the knowledge that *p*.

This is because—as we have found—some knowledge's having a stronger justificatory link to truth is not a circumstance that obtains once the justification is a *static link in place within the knowledge as a whole*. When suitably generalised, the redundancy problem shows, we saw, that a fallibly justified true belief that *p* is not a stronger static link to the truth that *p* than is the mere true belief that *p*. Correlatively, if there is to be a stronger link to the truth that *p* due to the presence of supportive justification, this needs instead to be an *active* justificatory linking to truth. In other words, its linking to the truth that *p* would be achieved *not* by the justification's being a part of the (completed) state that is the knowledge that *p*. Rather, we might treat an active linking along such lines as more of a preparatory link, one that could help us to *proceed* to reach the knowledge—yet without its thereby becoming a *part* of the knowledge.

That interpretive prospect coheres with Sect. 5's acknowledgement that the redundancy problem leaves untouched the uncontentious-because-trivial idea that infallible justification per se is stronger *as* justification than is fallible justification per se; and, of course, it is equally uncontentious-because-trivial that fallible justification per se is stronger *as* justification than is a complete absence of justification. Even the redundancy problem allows that, whenever you have infallible rather than fallible justificatory support for believing that *p*, you have justificatory support which—purely *as* justification per se—is stronger. But we should not infer that there is thereby a stronger link *to truth*. The redundancy problem implies that whenever such justification is

included within knowledge—so that now we are considering at least a justified true belief as a completed combination—this resulting combination does *not* possess such an increased strength as a link to truth. The redundancy problem enters the epistemological story when we are comparing instantiations of the property InFJT_B with ones of the property FJT_B, or (in this section) instantiations of FJT_B with ones of TB. However, what the redundancy problem does not impugn is the capacity for justification per se—let alone for stronger rather than weaker justification (for example, instantiations of InFJ rather than of FJ)—to lead a person, actively, *to* the state of having a true belief. The redundancy problem does not bar that *active* linking to truth from being more, or less, literally *effective* in a truth-linked way: it could well be stronger, or indeed less so, in that active way in accord with the justification's being more, or indeed less, strong.

All of that leads us constructively towards an unorthodox picture—which we may call knowledge-*minimalism*—that was advanced memorably by Sartwell (1991, 1992), in particular. He argued that knowledge *is* simply true belief: nothing more, nothing less. His picture has attracted few adherents, seemingly because it weakens the conceptual link between knowing that *p* and having justification for a true belief that *p* (that is, for the true belief that, according to knowledge-minimalism, is *ipso facto* her knowledge that *p*). But might that standard reaction—that wariness about knowledge-minimalism—be *needlessly* worried about knowledge-minimalism's prospect of doing justice to how knowledge and justification are to be linked within our conceptual theorizing? Even a minimalist conception of knowledge need not deny that knowledge does, or at least can, have *something* important to do with justification; what, though? Sartwell's answer (1991, p. 161) was that, always, the justification is 'a *criterion*, though not a logically necessary condition, of knowledge'. This is not the epistemologically traditional view. Is it at least coherent, though?

Indeed so, especially given this paper's argument. For knowledge-minimalism can be conjoined smoothly with one or another coherent way of recognising the epistemic contribution being made by some evidence, for instance, to one's knowing on a given occasion. One such form of conjunction opens the conceptual door to the idea of there being *different epistemic strengths* of knowledge, even for knowledge of a single *p*. On that gradualist approach (as it has been termed), any instance of knowledge that *p* is itself better or worse—possessing *some* epistemic grade—as knowledge that *p*. Goldman (1999, pp. 23–26) allows there to be weak knowledge, strong knowledge, and superstrong knowledge. For Goldman, weak knowledge is mere true belief (this being, for Sartwell, what *all* knowledge is). I have elsewhere (2001, Chap. 4) envisaged a more *extended* acceptance of possible strengths of knowledge that *p*, for a specific *p*. I allowed—as a limiting case for the concept of knowledge—the minimalist possibility that some knowledge that *p* might be merely a true belief that *p*. But then I allowed that further instances of knowledge that *p* could include increasingly strong justificatory support, with the knowledge that *p* being correlatively strengthened itself as knowledge that *p*—that is, being improved epistemically as knowledge of that specific *p*. In principle, then, I allowed that there can be both minimal knowledge that *p* and many possible grades of improved knowledge that *p*.¹⁹ Foley's (2012) view of knowledge is

¹⁹ These grades of potential improvement in one's knowledge that *p* are purely epistemic, in accord with any justificatory improvement within the knowledge. That sort of improvement must be distinguished

also apposite here. He conceives of knowledge (almost) purely as mere true belief. I say ‘(almost)’ because he conceives of knowledge as true belief plus enough important information (where information is also true belief); yet this further condition required by Foley could be viewed as a form of epistemic justification, a point noted by Warenski (2014, p. 896). And so perhaps he, too, could—although in fact he does not—talk in terms of there being minimal knowledge that *p* and improved knowledge that *p*.²⁰

Those views from Goldman, Foley, and me are not clearly *pure* forms of knowledge-minimalism, because each, it seems, allows that, for a given *p*, there could be knowledge that *p* that literally includes some justification. In contrast, Sartwell’s view—like this paper’s—is more *starkly* knowledge-minimalist. For, again, this view says that knowledge is *only ever*, in itself, a true belief—and hence that, even if there is justification for a given such instance of knowledge, this does not literally affect or change the epistemic nature of that knowledge. Instead, the knowledge, *as* knowledge, remains only the true belief; any justification for that belief’s being true—even though the justification is thereby *supporting* the knowledge—is not also literally a *part* of the knowledge. How, then, is this stark knowledge-minimalism to accommodate the fact that justification often *is* present with—indeed, because it has often generated—a particular instance of knowledge? We should ask whether all knowledge is only *barely* knowledge, according to this strict form of knowledge-minimalism.

And the answer to that pressing question would be that in one sense all knowledge *is* barely knowledge, even while in another sense it need not be so. *Yes*: on this paper’s knowledge-minimalism, all knowledge is barely knowledge, in the sense that something’s being knowledge *consists merely in* its being a true belief—with its being this true belief thus being considered independently of whatever, if any, further justification for it might also be present. On the other hand, *no*: an instance of knowledge need not barely be knowledge, in the sense that it might not have to *function epistemically on its own merely as* a true belief. For it could function as part of a larger epistemic ‘package’, one that at least often includes some justification supporting and supplementing the knowledge in the ways that justification does this. In that sense, a given instance of knowledge (a true belief) might actually be justified quite well, thereby enjoying all of the epistemic and psychological security that would typically attend the presence of that justification. But even this would not change the knowledge’s *being* knowledge only insofar as it is a true belief.

Footnote 19 continued

from mere *psychological* strengthening in believing that *p*, such as when one is hypnotised into holding a true belief that *p* more doggedly. After all, that psychological strengthening need not be accompanied by any epistemic strengthening; and so a belief that is held with more conviction is *not* thereby improved as knowledge that *p*, on that account of mine. Note, though, that my earlier approach was general enough to accommodate different conceptions of purely epistemic improvement. Congruently, for instance, Pritchard (forthcoming) argues that our assessments of the sort of luck or risk that—he also argues—most of us think is incompatible with knowing should be understood in modal rather than probabilistic terms, and that such terms are *gradational*, allowing a knowledge-precluding false belief, say, to be *more or less* modally close to our actual world.

²⁰ Lycan (1994) and Kvanvig (2003) argue against the cogency of Sartwell’s position; and I would expect them to apply their arguments likewise against these further forms of knowledge-minimalism. For criticism of their arguments, however, along with some experimental support for at least part of Sartwell’s position, see Sackris and Beebe (2014).

Accordingly, here is how we could synthesize the previous paragraph's combination (that 'in one sense all knowledge *is* barely knowledge, even while in another sense it need not be so'):

Any instance of knowledge is simply a true belief: that is, the belief's *being* knowledge is nothing beyond its being true. Even so, in principle an instance of knowledge could be justified, to some or another extent; maybe all or most knowledge will in fact be justified. Hence, in principle there can be justified knowledge, perhaps a lot of it. But my main point has been that to describe an instance of knowledge as being justified is not to speak empty, because in principle any instance of knowledge could also *fail* to be justified. So, even when—as might in fact always happen—an instance of knowledge *is* justified, this does not entail that the true belief in question would not have been knowledge *until* it was justified. Any instance of justified knowledge is thereby a justified true belief—but not because the knowledge *in itself* is at least a justified true belief.

Even on knowledge-minimalism, therefore, justification can continue to be accorded epistemic roles that accord with its traditionally being so closely associated with knowing. The redundancy problem described in this paper—along with our distinction between two ways of being linked by justification to truth—provides a conceptual framework within which we may usefully regard the basic idea behind knowledge-minimalism more favourably in that respect than most epistemologists have done. Again, as I indicated just now, we may continue to expect that justification might well in fact play one or more roles in how we know. But the moral that we should hold in mind is that the main such role would remain a metaphysically *extrinsic* role—a role that is not metaphysically *constitutive* of knowing. The contribution would instead be one of being merely *causally* constitutive of some knowing's coming to exist.²¹ That causal contribution might admit of grades, too: we will at least hope that (with all else being equal) better justification is better at *giving* us knowledge, at *getting* us knowledge.

I am proposing, then, that the redundancy problem supports the following picture:

Whenever a person has a true belief that *p*, this is knowledge that *p*. If the person also has (and has used aptly)²² some pertinent justificatory support for *p*'s being

²¹ The justification's being possessed could also be causally constitutive of mere *attempts*—including unsuccessful ones—to bring into existence some instance of knowing. And of course the causal constitutiveness of which I am talking need not be the *only* causal contributor to the existence of the attempt or the knowing. In that respect, then, I am simplifying the story. My immediate aim is to note simply the broad categorial point—namely, that we may wish to regard justification as a part of our epistemic lives (including our coming to know) that is never literally a part of the knowledge produced at least partly via the justification. (Some justification might itself include knowledge. That is a distinct point.) Theory-of-knowledge would thus be conceptually distinct from theory-of-justification—a distinction long embraced by Foley (1987).

²² Many epistemologists have remarked on the need for evidence to be used aptly, not merely to be present (even if thereby *able* to be used aptly). For example, Kornblith (1980) noted as much on behalf of a (reliabilist) grounding condition. I have described in detail (2013) the importance of an 'activist' conception of evidence within sceptical arguments.

the case, this justification has, or at least could have, been what *brought* the person to, or at least towards, the state of having that knowledge.

That would be an *active* justificatory linking of the person to a true belief—to the knowledge. We may then distinguish between at least these two kinds of way in which such a link might be effected:

Internalist The linking could be a process of self-aware guidance, such as by involving the deliberate gathering and assessment of evidence prior to forming what one hopes will be a true belief, indeed knowledge. (This would be an internalist paradigm of an active justificatory genesis for the true belief.)

Externalist The justificatory process might be the instantiation of a reliable belief-forming process, with the reliability as such being metaphysically constitutive of the justification's presence. (This would be an externalist paradigm of an active justificatory genesis for the true belief.)²³

In either of those kinds of case, the following conceptual option emerges for us as epistemological interpreters:

In general, some justification might have been causally constitutive (in either an internalist or an externalist way) of some knowledge's coming to exist—without the justification's thereby being a *part* of the resulting knowledge. In any given case of knowing, perhaps that knowledge would not in fact have come to exist, if not for the justification's existing and functioning aptly. Yet even this would not entail the justification's being literally a part of the knowledge that would have now (at least partly by way of the active use of the justification) come to exist.

Nothing in that picture is at odds with the possibility that in fact all instances of our knowledge are accompanied, or even generated, by some justification. Whether that does occur depends on what justification is, for a start; and what justification is constrains also how we would try to ascertain whether in fact all of our knowledge is accompanied or generated by justification. (Maybe it is an empirical matter; maybe it is a transcendental need; etc.) Regardless, however, of what the outcome would be of any such investigation (be it empirical or be it otherwise), my basic conceptual point remains:

If we wish to view a person's justification for believing that *p* as strengthening her link as a knower to the truth that *p*, we would do well to regard that justification as not literally being *part* of that (completed state of) knowledge. We would do better to regard that justification as instead a causal precursor to the knowledge as such—so that the justification's role is that of actively linking the person to the truth that *p* by helping to *produce* the true belief that *p* that is *ipso facto* her knowledge that *p*.

The redundancy problem has thus opened the door to this alternative conceptual prospect within the metaphysics of knowledge.

²³ This conception of the distinction between epistemic internalism and epistemic externalism was introduced above, in note 4.

And that prospect deserves a name of its own. I have been calling it knowledge-minimalism. And I have explained how it is the thesis that knowledge as such is merely a true belief. Once more, I stress, this is not thereby a complete rejection of the epistemic significance of epistemic justification. Knowledge-minimalism is compatible with such justification's having many forms of epistemic significance, such as by being a guide at any given moment to *further* knowledge—that is, to knowledge not as yet acquired. It is also compatible with all knowledge's in fact being productively supported by justification, so that wherever and whenever there is knowledge there either is or has been justification, playing a relevantly active role in the knowledge's coming to be. But knowledge-minimalism distinguishes between that form of relationship (even when the relationship does obtain) and knowledge's ever needing to *include within* itself that causally associated justification. We saw that good evidence or strong reliability, say, can help you to know that p—by being good *sources*, in their respective ways and with all else being equal, of the actively generated true belief that p—without their having to be literally a *part* of the resulting knowledge that p.

Maybe our lives will never actually include some knowledge that is completely unsupported by justification: maybe any knowledge that in fact we will ever have will be supported by justification that we will also have. Even if that is what transpires, however, there will remain a need for epistemologists not to confuse such a *conjunction* of circumstances or, more strongly, such a circumstantial *progression* with a *metaphysically constitutive inclusion*—by inferring that therefore the justification is a part, rather than merely a helpful generator, of the knowledge. And the redundancy problem, I suggest, should make us more alert to this possible conceptual refinement of some standard epistemological thinking. We have this opportunity to travel far—even further than we might well have set out to travel—from a Cartesian knowledge-infallibilism.

Acknowledgements I am grateful to Dick Foley, Brent Madison, and two anonymous referees for comments on some earlier drafts of this paper.

References

- Bates, J. (2013). Damning the swamping problem, reliably. *Dialectica*, 67, 103–116.
- Chalmers, D. C. (2002). Does conceivability entail possibility? In T. S. Gendler & J. Hawthorne (Eds.), *Conceivability and possibility* (pp. 145–200). Oxford: Oxford University Press.
- Clarke, D. M. (2012). The epistemology of Descartes. In S. Hetherington (Ed.), *Epistemology: The key thinkers* (pp. 90–110). London: Continuum.
- Davis, W. A., & Jäger, C. (2012). Reliabilism and the extra value of knowledge. *Philosophical Studies*, 157, 93–105.
- Dougherty, T. (2011). Fallibilism. In S. Bernecker & D. Pritchard (Eds.), *The Routledge companion to epistemology* (pp. 131–143). New York: Routledge.
- Dutant, J. (2013). In defence of swamping. *Thought*, 2, 357–366.
- Engel, M. (1992). Is epistemic luck compatible with knowledge? *The Southern Journal of Philosophy*, 30, 59–75.
- Engel, M. (2011). Epistemic luck. In *The Internet Encyclopedia of Philosophy*. <http://www.iep.utm.edu/epi-luck>.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. New York: Oxford University Press.
- Foley, R. (1987). *The theory of epistemic rationality*. Cambridge, MA: Harvard University Press.

- Foley, R. (2012). *When Is true belief knowledge?*. Princeton, NJ: Princeton University Press.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23, 121–123.
- Goldman, A. I. (1999). *Knowledge in a social world*. Oxford: Clarendon Press.
- Hetherington, S. (1996). *Knowledge puzzles: An introduction to epistemology*. Boulder, CO: Westview Press.
- Hetherington, S. (1999). Knowing failably. *The Journal of Philosophy*, 96, 565–587.
- Hetherington, S. (2001). *Good knowledge, bad knowledge: On two dogmas of epistemology*. Oxford: Clarendon Press.
- Hetherington, S. (2005). Fallibilism. In *The Internet Encyclopedia of Philosophy*. <http://www.iep.utm.edu/f/fallibil.htm>.
- Hetherington, S. (2011). The Gettier problem. In S. Bernecker & D. Pritchard (Eds.), *The Routledge companion to epistemology* (pp. 119–130). New York: Routledge.
- Hetherington, S. (2013). Skeptical challenges and knowing actions. *Philosophical Issues*, 23, 18–39.
- Hetherington, S. (2016). *Knowledge and the Gettier problem*. Cambridge: Cambridge University Press.
- Hetherington, S. (forthcoming a). Understanding fallible warrant and fallible knowledge: Three proposals. *Pacific Philosophical Quarterly*.
- Hetherington, S. (forthcoming b). Skepticism and fallibilism. In D. Machuca & B. Reed (Eds.), *Skepticism: From antiquity to the present*. London: Bloomsbury.
- Hintikka, J. (1962). *Knowledge and belief: An introduction to the logic of the two notions*. Ithaca, NY: Cornell University Press.
- Howard-Snyder, D., Howard-Snyder, F., & Feit, N. (2003). Infallibilism and Gettier's legacy. *Philosophy and Phenomenological Research*, 66, 304–327.
- Kornblith, H. (1980). Beyond foundationalism and the coherence theory. *The Journal of Philosophy*, 77, 597–612.
- Kvanvig, J. L. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge: Cambridge University Press.
- Lycan, W. G. (1994). Sartwell's minimalist analysis of knowing. *Philosophical Studies*, 73, 1–3.
- McDowell, J. (2011). *Perception as a capacity for knowledge*. Milwaukee, WI: Marquette University Press.
- Olsson, E. J. (2011). Reply to Kvanvig on the swamping problem. *Social Epistemology*, 25, 173–182.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Clarendon Press.
- Pritchard, D. (2007). Anti-luck epistemology. *Synthese*, 158, 277–297.
- Pritchard, D. (2009). *Knowledge*. Basingstoke: Palgrave Macmillan.
- Pritchard, D. (2013). Knowledge cannot be lucky. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary debates in epistemology* (2nd ed., pp. 152–164). Malden, MA: Wiley Blackwell.
- Pritchard, D. (forthcoming). Epistemic risk. *The Journal of Philosophy*.
- Reed, B. (2000). Accidental truth and accidental justification. *The Philosophical Quarterly*, 50, 57–67.
- Reed, B. (2002). How to think about fallibilism. *Philosophical Studies*, 107, 143–157.
- Reed, B. (2012). Fallibilism. *Philosophy Compass*, 7(9), 585–596.
- Riggs, W. (2002). Reliability and the value of knowledge. *Philosophy and Phenomenological Research*, 64, 79–96.
- Sackris, D., & Beebe, J. R. (2014). Is justification necessary for knowledge? In J. R. Beebe (Ed.), *Advances in experimental epistemology* (pp. 175–192). London: Bloomsbury.
- Sartwell, C. (1991). Knowledge is merely true belief. *American Philosophical Quarterly*, 28, 157–165.
- Sartwell, C. (1992). Why knowledge is merely true belief. *The Journal of Philosophy*, 89, 167–180.
- Shope, R. K. (1983). *The analysis of knowing: A decade of research*. Princeton, NJ: Princeton University Press.
- Sosa, E. (1999). How must knowledge be modally related to what is known? *Philosophical Topics*, 26, 373–384.
- Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. I). Oxford: Clarendon Press.
- Sosa, E. (2011). *Knowing full well*. Princeton, NJ: Princeton University Press.
- Warenski, L. (2014). Review of Foley (2012). *Mind*, 123, 894–898.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Clarendon Press.
- Zagzebski, L. (1994). The inescapability of Gettier problems. *The Philosophical Quarterly*, 44, 65–73.
- Zagzebski, L. (2003). The search for the source of epistemic good. *Metaphilosophy*, 34, 12–28.