

Revisionary Epistemology

Author 1 *Davide Fassio, University of Geneva*
davide.fassio@unige.ch

Author 2 *Robin McKenna, University of Vienna*
rnmckenna@gmail.com

Abstract What is knowledge? What should knowledge be like? Call an epistemological project that sets out to answer the first question ‘descriptive’ and a project that sets out to answer the second question ‘normative’. If the answers to these two questions don’t coincide - if what knowledge should be like differs from what knowledge is like - there is room for a third project we call ‘revisionary’. A revisionary project starts by arguing that what knowledge should be differs from what knowledge is. It then proposes that we revise our account of knowledge accordingly. Our aim in this paper is to develop a methodology for revisionary projects in epistemology. Put roughly, the thought is that we start by looking at the various things that we expect knowledge to do for us. Once we have a list of the various things we expect knowledge to do for us we have a ‘job description’; a list of tasks we need done, and that we expect knowledge to perform. With the job description in hand, we can ask what knowledge would have to be like in order to perform these tasks. Along the way we give some reasons for embarking on a revisionary project in epistemology, and we outline what the upshot might be.

0. Introduction

What is knowledge? What should knowledge be like? Call an epistemological project that sets out to answer the first question ‘descriptive’ and a project that sets out to answer the second question ‘normative’. If the answers to these two questions don’t coincide - if what knowledge should be like differs from what knowledge is like - there is room for a third project we call ‘revisionary’. A revisionary project starts by arguing that what knowledge should be differs from what knowledge is. It then proposes that we revise our account of knowledge accordingly, i.e. in accordance with what knowledge should be like.

We will say more about the relationship between these projects, but from the outset let us make two things clear. First, when we talk about what ‘knowledge’ is or should be we are primarily talking about the concept knowledge. To say that knowledge is (should be) a certain way is to say that the concept is (should be) a certain way. Space for a revisionary project opens up if the concept isn’t as it should be. This is not to say one can’t make sense of our three projects at the level of properties. In terms of properties, the descriptive project is interested in the referents

picked out by our knowledge concept, whereas the normative project is interested in the referents that the concept should be picking out. Again, space for a revisionary project opens up if the concept isn't picking out the referents it should pick out. But, given our interest in revisionary projects, we prefer to talk about the knowledge concept, not the property (we usually can't revise properties).¹

Second, we draw the distinction between the three projects in terms of the question they address, not in terms of the methodology used to address it. As will become clear, we think that the appropriate methodologies for descriptive, normative and revisionary projects overlap, at least to some extent. This has the practical consequence that it is sometimes unclear which of our three projects are being pursued by particular epistemologists (see §1).

While descriptive projects are common in contemporary epistemology, revisionary projects have been neglected. This neglect isn't particular to epistemology. In general, philosophers are more interested in descriptive projects than normative or revisionary projects (we discuss some exceptions in §1).² While there are many reasons for this, three are worth mentioning. First, absent any reason to think that what knowledge should be differs from what knowledge is, there is no motivation for revisionary projects. Second, revisionary projects involve an objectionable semantic or conceptual revisionism. Put bluntly, philosophers shouldn't be in the business of changing the meanings of our words. Third, it is unclear how one determines what knowledge should be like. In contrast, while epistemologists disagree about what knowledge is, most agree on the appropriate methodology for tackling the question. It involves a mixture of intuitions about hypothetical cases, general intuitions about knowledge (e.g. that it is valuable), and so on.

Our primary aim in this paper is to explain why these lines of resistance to revisionary projects are mistaken. In response to the first reason, we present evidence that our concept knowledge is governed by a set of inconsistent principles. On the assumption that knowledge shouldn't be inconsistent, it follows we have motivation to pursue a revisionary project in epistemology. In response to the second, we argue that, while philosophers perhaps shouldn't be in the business

¹ We can revise some properties (e.g. socially constructed properties, like the property of being rich), but the knowledge property isn't usually regarded as socially constructed (but see Kusch, *Knowledge by Agreement*). Because we want to set this issue aside, we'll not talk much about the knowledge property.

² To anticipate: In §1 we point out that revisionary projects were far more common in the first half of the 20th century, and suggest that the reason has to do with the perceived aims of philosophy. It is worth noting that there has recently been renewed interest in revisionary projects; see Cappelen, 'Illusions of Thought'; Chalmers, 'Verbal Disputes'; Haslanger, 'What Knowledge Is', 'What Good Are Our Intuitions?'

of stipulating new meanings for our words, there is no reason why philosophers shouldn't be in the business of suggesting improvements.

In response to the third, we outline how one might pursue a revisionary project in epistemology. Put roughly, to figure out what knowledge should be like we need to look at the various things we expect knowledge to do for us. For instance, it is frequently said that knowledge has the function of identifying good informants. In saying that Sarah knows the train leaves at 3pm I identify her as a good informant on the matter of when the train leaves.³ Once we have a list of the various things we expect knowledge to do for us we have a 'job description'; a list of tasks that we need done, and that knowledge is supposed to perform.⁴ With the job description in hand, we can ask what knowledge would have to be like in order to perform these tasks. In answering that question we figure out what knowledge should be like.

Our secondary aim is to illustrate how our revisionary methodology works. To do this we will take a number of functions that we expect knowledge to serve and show that the upshot is a view we call 'revisionary sensitive invariantism'. The view is 'revisionary' because it is an account of what knowledge should be, not what knowledge is. The view is 'sensitive invariantist' because it says that whether a subject *S* knows some proposition *p* depends on various practical factors, as well as the usual truth-conducive factors.⁵ While we are inclined to think that revisionary sensitive invariantism is the best revisionary view available, our purpose here is mainly illustrative. If we can convince you that the methodology is sound we will have achieved our goal.

We proceed as follows. In §1 we situate our paper with respect to its historical and contemporary precedents. In §2 we present evidence that our concept knowledge is governed by inconsistent principles, and we argue that this evidence motivates a revisionary project in epistemology. In §3 we outline and illustrate our revisionary methodology.

1. Some Background

In this section we briefly survey historical and contemporary revisionary approaches. We close the section by offering a tentative diagnosis and endorsement of the driving thought behind revisionary approaches, both within and outside epistemology.

³ Craig, *Knowledge and the State of Nature*, 11-17

⁴ We borrow the 'job description' metaphor from Fisher, 'Meanings and Methodologies', 60-1.

⁵ Fantl and McGrath, *Knowledge in an Uncertain World*; Hawthorne, *Knowledge and Lotteries*; Stanley, *Knowledge and Practical Interests*.

We can start with the method of explication, which one finds in both Carnap and Quine. Put roughly, the idea behind explication is that one takes an ordinary language concept (the ‘explicandum’) and tries to find another concept (the ‘explicatum’) which fulfils the following criteria:

- i. The explicatum is similar to the explicandum in that, in many (not necessarily all) cases in which the explicandum applies, the explicatum does too.
- ii. The characterisation of the explicatum is exact.
- iii. The explicatum is fruitful (e.g. for the formulation of universal statements).
- iv. The explicatum is simple.⁶

While more could be said about explication, what is important here is that Carnap’s criteria capture both the conservative and revisionary aspects of explication. On the one hand, it is a requirement that the explicatum be similar to the explicandum, in that their extensions must not differ too much. On the other, differences are permitted, and given criteria (ii)-(iv) they are to be expected. As Quine put it, in offering an explication of some concept:

We do not claim synonymy. We do not claim to make clear and explicit what the users of the unclear expression had unconsciously in mind all along. We do not expose hidden meanings, as the words ‘analysis’ and ‘explication’ would suggest; we supply lacks. We fix on the particular functions of the unclear expression that make it worth troubling about, and then devise a substitute, clear and couched in terms to our liking, that fills those functions.⁷

In short, the aim in explicating a concept is to improve it. But there is a difference between improving a concept and inventing a new one, and both Carnap and Quine saw explication as a matter of improvement, not invention.

Turning to contemporary epistemology, Edward Craig has provided a ‘practical explication’ of knowledge, by which he means an account of what knowledge would have to be like in order to play a particular functional role (identifying good informants).⁸ It is no part of Craig’s view that his account of what knowledge would have to be is a perfect match for the actual concept, although he thinks it is important that the differences be minor. Here’s Craig:

We take some prima facie plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like, what

⁶ Carnap, *Foundations of Probability*, 5-8.

⁷ Quine, *Word and Object*, 258-9.

⁸ Craig, *Knowledge and the State of Nature*.

conditions would govern its application. Such an investigation would still have an anchorage point in the everyday concept: should it reach a result quite different from the intuitive intension, or one that yielded an extension quite different from the intuitive extension, then, barring some special and especially plausible explanation of the mismatch, the original hypothesis about the role that the concept plays in our life would of course be the first casualty. For it is not the idea to construct an imaginary concept, but to illuminate the one we actually have, though it be vague or even inconsistent.⁹

More recently, Sally Haslanger has provided an account of what knowledge would have to be like in order to further certain social and political goals.¹⁰ However, Haslanger is more open to revision than Craig:

[T]he specifically epistemic questions "What is knowledge?" or "What is objectivity?" require us to consider what work we want these concepts to do for us; why do we need them at all? The responsibility is both to investigate our purposes in having them, and then to define them in a way that best meets our legitimate purposes. In doing so we will want to be responsive to ordinary usage (and to aspects of both the connotation and extension of the terms). However, there is also a stipulative element to the project: this is the phenomenon we need to think about; let us use the term 'knowledge' to refer to it.¹¹

Craig, Haslanger and Carnap differ with respect to how much emphasis they put on Carnap's first criterion for successful explication (Craig puts a lot of emphasis on it whereas Carnap and Haslanger a lot less) and the sorts of roles that are relevant (Craig and Haslanger are interested in practical roles whereas Carnap is interested in theoretical roles). But, despite these differences, they share the core idea that, when developing an account of a concept such as knowledge, one should do so with an eye to the purpose or roles of that concept.

It is an interesting sociological question why, with exceptions like Craig and Haslanger, explicitly revisionary approaches have gone out of favour (we leave it up to the reader to decide how common implicitly revisionary approaches are).¹² Our (very) tentative diagnosis is that it has

⁹ Ibid., 2

¹⁰ Haslanger, 'What Knowledge Is'.

¹¹ Ibid., 468.

¹² Consider the project of providing an informative 'analysis' of the concept knowledge, which is often seen as the paradigm of an attempt to fit an account of a concept around a set of intuitions. But, nowadays anyway, considerations about the value and functions of knowledge are central to the project, and its practitioners are open to revision. For instance, we might have reason to reject an account that deals with all the cases on the grounds that it fails to account for the value or

something to do with the perceived aims of philosophy. When you read Carnap it is clear that he thinks of philosophical analysis as a tool. We can use the methods, argumentative devices and vocabulary of philosophy to develop sophisticated accounts of concepts, such as that of probability, which can then be put to work in other disciplines (in this case, science). As Herman Cappelen has put it:

These authors' [Carnap, Quine, and others] interest in describing how language worked served, in large part, a normative goal: to figure out how to make language better and how to improve it as a cognitive tool.¹³

The underlying idea is that the ultimate aims of philosophy are ameliorative, not descriptive. One can't entirely separate the ameliorative aim from the descriptive. It is hard to improve something, such as language, without first finding out how it works. But, for someone like Carnap, finding out how language works is the start of the task, not the goal.

We suspect that, unlike Carnap and Quine, most philosophers in the last fifty or so years don't think the ultimate aims of philosophy are ameliorative.¹⁴ While we lack the space to argue the point here, we think that Carnap and Quine had the right view about the ultimate aims of philosophy. Much (though not all) philosophy is concerned with how we think about the world around us, and with 'analysing' concepts - figuring out their intensions, extensions, and so on. But concepts are tools for making sense of the world, and expressing truths about what is going on it. If it turns out that we are better able to talk about the world by changing our concepts - by modifying and improving our tools - then we have reason to do so. Applying this to knowledge, our concept knowledge is a tool we use to identify a particular sort of cognitive success. But it could be that our current concept isn't ideally suited to identifying the relevant sort of success. Unlike Carnap and Quine we don't focus on the role that knowledge might play in a formal system (but see §3.3). Rather, like Craig and Haslanger, we will look at the practical roles that we expect knowledge to play, and develop an account of what knowledge would have to be like in order to play those roles. The end result is an improvement on our current concept.

functions of knowledge (see Jones, 'Do We Value Knowledge?'; Pritchard, 'On Meta-Epistemology'; Riggs, 'The Value Turn').

¹³ Cappelen, 'Nonsense and Illusions of Thought', 23.

¹⁴ Cappelen provides an impressive list including philosophers of language like Davidson, Kripke and Putnam and linguists like Heim and Kratzer (ibid.).

2. Knowledge and Inconsistency

The argument of this section is as follows:

1. If knowledge isn't as it should be, we have reason to embark on a revisionary project.
2. Knowledge isn't as it should be.
3. We have reason to embark on a revisionary project.

We take the first premise to require little by way of defence. As we put it earlier, if knowledge isn't as it should be, space opens up for a revisionary project.

Our defence of the second premise will start with evidence that our concept knowledge is governed by a set of inconsistent principles (§2.1). There are a number of things one could say about this evidence, but two are particularly important. First, one might say that, while all of these principles seem to govern knowledge, they don't all actually do so. Second, one might say that our concept knowledge is actually inconsistent. In §2.2 we argue against the first response, and in favour of the second.¹⁵ While we take our argument to only make a *prima facie* case that knowledge is inconsistent, the *prima facie* case should be more than enough to establish the second premise. That the case can be made at all is excellent evidence that the concept knowledge is 'messy', in the sense that our best thinking about it is pulled in conflicting directions. Paraphrasing Craig, in a revisionary project we revise a concept we actually have *because* it is messy, or even incoherent.

The argument requires the further premise that, if knowledge is inconsistent, it isn't as it should be. It is very plausible that, in general, our concepts shouldn't be inconsistent and that, for any concept, if it turns out that that concept is inconsistent, we should do something about it, viz. change it in such a way that it is no longer inconsistent. But Matt Weiner has argued that, while knowledge is inconsistent, the inconsistency is (mostly) harmless.¹⁶ While we lack the space to go into any detail, we doubt that the inconsistency is harmless. Weiner's basic thought is that the principles governing knowledge only give inconsistent verdicts in marginal cases.¹⁷ But, as will become clear below, this isn't so. This is partly because we focus on a wider range of principles than Weiner, and so uncover a lot more inconsistency. But it is also because the principles give conflicting verdicts in very mundane cases. For instance, some of our principles tell us that I can

¹⁵ For other arguments to this effect see Schiffer, 'Contextualist Solutions'; Weiner, 'Inconsistency of Knowledge Ascriptions'. While our argument owes much to Weiner, there is an important difference; see immediately below.

¹⁶ Weiner, 'Inconsistency of Knowledge Ascriptions'.

¹⁷ *Ibid.*, 14-18.

know a train time just from being told by a reliable passerby, whereas other principles tell us that I can't. If the principles give conflicting verdicts even in mundane cases, we can expect this to cause problems.

2.1. The Evidence

We will introduce the principles in three groups. First, we have (Parity of Evidence Principle) and (Knowledge-Action Principle):

(Parity of Evidence Principle) If two subjects S and S^* have the same evidence for some proposition p , S is in a position to know p iff S^* is in a position to know p .¹⁸

(Knowledge-Action Principle) It is rational for S to treat p as a reason for action iff S knows p .

(Parity of Evidence Principle) fits with the epistemological orthodoxy that whether a true belief counts as knowledge just depends on whether the evidence supporting it is sufficiently strong. According to this orthodoxy, differences in practical factors - e.g. what's at stake - are irrelevant. The principle is also supported by two intuitions. First, sentences like (1) seem absurd (where it is specified that S hasn't gained or lost any evidence):

(1) S knows p now, but he didn't yesterday, when it mattered more.¹⁹

Second, knowledge is stable: it does not come and go with variations in interests, or other factors unrelated to truth.²⁰

(Knowledge-Action Principle) is supported by ordinary evaluations of the rationality of actions.²¹ For instance, a doctor is criticisable for using a needle he does not know is safe, and Sally is criticisable for not paying her health insurance as she does not know she won't fall ill. Similarly, if I turn left on the way to the airport I can justify myself by saying I know the airport is in that direction.

However, (Parity of Evidence Principle) and (Knowledge-Action Principle) seem to be inconsistent. Whether it is rational to treat a proposition as a reason for acting is sensitive to practical factors. Jane has to catch a train, but she doesn't know which platform it leaves from. Suppose Jane is in a low stakes situation and she asks a reliable looking passerby which platform

¹⁸ We're using 'evidence' in a very broad sense here that is intended to cover all and only those facts that are conducive to p being true.

¹⁹ DeRose, 'Now You Know It'; Yourgrau, 'Knowledge and Relevant Alternatives'.

²⁰ This point goes back to Plato, but see Williamson, *Knowledge and its Limits*, 78-80.

²¹ Hawthorne and Stanley, 'Knowledge and Action'.

it is. Trusting her advice, she boards the relevant train. Intuitively, in doing so she is rational. Now suppose Jane is in a high stakes situation (if she misses her meeting she loses her job) and she asks the same passerby. In this case she would be irrational in boarding the train without seeking corroboration (e.g. asking an official). Because Jane would be rational in the first case but not in the second, it follows from (Knowledge-Action Principle) that she knows in the first case but not the second. However, because her evidence is exactly the same in both cases, it follows from (Parity of Evidence Principle) that she knows in the first case iff she knows in the second. So the two principles are inconsistent.

Second, we have (Testimony Principle) and (Social Principle):

(Testimony Principle) If S knows p , then S^* would be in a position to know that p if S were to tell her so (assuming S is generally trustworthy and nothing else interferes with the transmission of knowledge).

(Social Principle) Whether it is right to ascribe knowledge to a subject S can vary with the social context.

Many of our beliefs about the world around us come from testimony. So, if we want to avoid scepticism, something like (Testimony Principle) must be true. However, (Testimony Principle) is also incompatible with (Knowledge-Action Principle). According to (Testimony Principle), Jane is in a position to know based on the passerby's testimony in both the low and high stakes cases whereas, according to (Knowledge-Action Principle), she isn't.

(Social Principle) is supported by a variant of the Jane case. Say the passerby, Tom, knows which platform the train leaves from on the basis of further testimony from an official. If Tom has no reason to think a lot rides on Jane getting the right train, it is right for him to tell her that he knows which train it is. In contrast, if Tom has some reason to think a lot rides on Jane getting the right train, it is wrong for him to tell her that he knows which train it is.

However, (Social Principle) seems to conflict with (Parity of Evidence Principle) and (Testimony Principle). According to (Parity of Evidence Principle), whether a subject knows is just a matter of the evidence she possesses, and has nothing to do with the social context. According to (Testimony Principle), if the official who tells Tom which train the platform leaves from knows,

then Tom does too. But, in the case where Tom thinks the stakes are high for Jane, he can't appropriately say that he knows.²²

Finally, we have these four principles:

(Odds Principle) Even if p is massively against the odds still, no matter how unlikely it is that p , one can't know not- p just based on knowing the odds.

(Fallibilist Knowledge Principle) Some knowledge is based on less than conclusive evidence.

(Strength of Evidence Principle) Where p and q are both true, if S's belief p is justified to degree E , S*'s belief q is justified to degree E^* , and E is much higher than E^* , then if S is not in a position to know p , S* is not in a position to know q .

(Closure Principle) If S knows p , and S knows that p entails q , and S comes to believe q solely on the basis of competently deducing it from p while retaining knowledge that p throughout the competent deduction, then S knows q .

To see the motivation for (Odds Principle), consider an example. No matter how unlikely it is that my lottery ticket will win, I can't know it won't win just based on knowing the odds. Suppose Sally discards her lottery ticket before the draw, and justifies herself by saying that the ticket will lose. Sally is subject to criticism, and the natural way of voicing the criticism is to say: "you didn't know that your ticket would lose."

(Fallibilist Knowledge Principle) is supported by mundane examples of knowledge. If I read in the newspaper that Manchester United won 2-1 last night then I don't have conclusive evidence that Manchester United won 2-1 - the newspaper could be unreliable. Still, unless I have reason to think the newspaper is unreliable, I know Manchester United won 2-1. Or consider this case, from Jonathan Vogel.²³ If I parked my car outside then I don't have conclusive evidence that it is still outside - there could be car thieves about. Still, unless I have reason to think there are car thieves about, I know my car is outside.

²² We're assuming it is right for Tom to self-ascribe knowledge iff he knows. One might object that this is a matter of the pragmatics, not semantics, of knowledge ascriptions (see Brown, 'Contextualism and Warranted Assertibility'; Gerken, 'Roles of Knowledge'; Rysiew, 'Context-Sensitivity of Knowledge'). While we can't address this issue here note, first, that the pragmatic interpretation has its own problems (see Blome-Tillmann, 'Knowledge and Implicatures') and, second, (Social Principle) isn't essential to our aims in this section, or the paper in general.

²³ Vogel, 'Counterexamples to Closure'.

(Strength of Evidence Principle) is supported by the same epistemological orthodoxy as (Parity of Evidence Principle). But it is significantly weaker. (Strength of Evidence Principle) just says that, where p and q are both true, if subject S has far stronger evidence in favour of p than subject S^* has in favour of q , then it can't be that S^* is in a position to know q but S isn't in a position to know p . If Ailsa has far stronger evidence that she has hands than Catriona has that the train leaves at 3pm, then it can't be that Catriona is in a position to know the train leaves at 3pm but Ailsa isn't in a position to know she has hands.

Finally, to see the motivation for (Closure Principle) consider the mathematician who proves one result, thereby gaining some knowledge, and then proves that her first result entails a second result. The mathematician thereby gains further knowledge.

However, (Odds Principle), (Weak Knowledge Principle) and (Strength of Evidence Principle) are inconsistent.²⁴ I get far stronger evidence that my lottery ticket lost from knowing the odds than I get that Manchester United won 2-1 from reading the newspaper (it is far more likely that the newspaper got the result wrong than it is that I won the lottery).²⁵ By (Strength of Evidence Principle), if I don't know my lottery ticket lost, I don't know Manchester United won 2-1. But by (Odds Principle) I don't know my lottery ticket lost, yet by (Weak Knowledge Principle) I do know Manchester United won 2-1.

Further, (Odds Principle), (Weak Knowledge Principle) and (Closure Principle) are inconsistent.²⁶ Consider this example.²⁷ A friend suggests I go on safari, but I tell her I can't afford to because my financial situation is modest. However, I have a ticket in a lottery with a big cash prize. The winning ticket has been drawn but has not been announced. I know that if I can't afford to go on safari then I haven't won the lottery. (Weak Knowledge Principle) tells us that I know that I can't afford to go on safari. However, (Odds Principle) tells us that I don't know I haven't won the lottery. So, by contraposing (Closure Principle), it follows that I don't know I can't afford to go on safari, contradicting (Weak Knowledge Principle).

²⁴ Ibid.

²⁵ We're assuming that strength of evidential support is tied to likelihood here. Smith, 'What Else Justification Could Be', presents an interesting challenge to this way of thinking about evidential support, but we lack the space to deal with it. Suffice it to say that the view we assume is the standard one.

²⁶ Vogel, 'Counterexamples to Closure'; Hawthorne, *Knowledge and Lotteries*, Chap. 1.

²⁷ Hawthorne, *Knowledge and Lotteries*, 1-3.

2.2. The Explanations

The obvious thing to say about our principles is that, while each principle is *prima facie* plausible, knowledge is actually governed by a consistent subset of the various principles. Making good on this idea requires, first, identifying which principles do and which principles don't govern knowledge and, second, giving an error-theoretic explanation why we mistakenly think that certain principles govern knowledge. Our aim in the rest of this sub-section is to suggest that this response won't work. Due to limitations of space we can't discuss the various ways of filling in the details. What we'll do instead is focus on one pair of principles, (Parity of Evidence Principle) and (Knowledge-Action Principle), and a particular response, on which the former should be accepted and the latter rejected.²⁸ The reasons the particular response won't work will hopefully generalise to other responses.

The simplest and most effective way to show that (Knowledge-Action Principle) should be rejected would be to show that there are cases in which the subject hasn't got knowledge, yet intuitively can act.²⁹ Critics of (Knowledge-Action Principle) have proposed that Gettier cases are of this type.³⁰ Smith has the justified true belief that someone in his office will get the job, but he doesn't know this because he's wrong about the identity of the person (he mistakenly thinks it will be Jones). Yet, one might think, Smith is rational in acting on his belief that someone in his office will get the job. If that's right, then (Knowledge-Action Principle) should be rejected.

Unfortunately, these cases haven't moved those who accept (Knowledge-Action Principle). A common response is that the putative counter-examples ignore the distinction between violating a norm and being to blame for violating that norm.³¹ The thought is that, while Smith violates (Knowledge-Action Principle), his violation is blameless (Smith wasn't aware he was violating the principle). However, this move doesn't advance the debate. It suggests that the force of the putative counter-example relies on ignoring the distinction between violating a norm and being to blame for violating that norm. But those who have proposed these counter-examples deny that the putative counter-example is a case of norm-violation, whether blameless or not.³²

This debate has reached a dialectical stalemate. Critics of (Knowledge-Action Principle) think it is obvious that there are cases in which a subject is rational in acting yet lacks knowledge (and

²⁸ For this response see Brown, 'Subject-Sensitive Invariantism'; Gerken, 'Warrant and Action'.

²⁹ Or that there are cases in which the subject has got knowledge, yet intuitively can't act. While we focus on one sort of case, what we say should generalise to this sort of case too.

³⁰ Brown, 'Subject-Sensitive Invariantism', 171-4; Gerken, 'Warrant and Action', 535-7.

³¹ Hawthorne and Stanley, 'Knowledge and Action', 572-4; Littlejohn, 'Unity of Reason', 139-40.

³² Gerken, 'Warrant and Action', 537-43, is explicit about this.

cases in which a subject has knowledge yet isn't rational in acting), whereas proponents of the principle think it isn't. We would suggest that, when a debate about an ordinary language notion such as knowledge reaches this sort of stalemate, this is *prima facie* evidence that the notion is inconsistent, and excellent evidence that the ordinary language notion is 'messy', in the sense that our best thinking about it is pulled in opposing directions. If this is right, then we have good reason to embark on a revisionary project. We now turn to how a revisionary project might look.

3. Revisionary Epistemology

How, exactly, do we go about revising knowledge? The 'revisionist' proposes devising a new concept, knowledge*, to replace our current concept, knowledge.³³ But what constraints are there on our choice of knowledge*?³⁴ There are a potentially infinite number of ways in which we could revise our concept, and the revisionist needs to tell us how to choose between them. In §3.1 we introduce the methodology we employ in choosing between different revisionary accounts. In §3.2 we illustrate our methodology by developing a revisionary account of knowledge* we call revisionary sensitive invariantism. We finish in §3.3 by briefly mentioning a further benefit of revisionary sensitive invariantism.

3.1. Our Revisionist Methodology

We can start by distinguishing between 'free' and 'constrained' revision. As examples of free revision, consider stipulating meanings and using words in a 'technical sense'. While free revision is perfectly fine in some contexts, we doubt there is much to be gained by developing a freely revisionary account of knowledge*. In general, when the target of one's analysis is a concept X, one may introduce stipulated meanings for terms that feature in the analysis of X, but not for X. Since the target of our analysis is the concept knowledge, we can't just stipulate a meaning for knowledge*. This leaves us with constrained revision, which preserves certain features of the

³³ We use the locution 'revise the concept' to communicate the intuitive idea that we can revise and improve our concepts. How this intuitive idea gets cashed out depends on your view of concepts. If concepts are unchanging entities then the revisionist is proposing taking a word ("knowledge") and using it to express a different but related concept. If you prefer to think of concepts as closer to cognitive tools - tools for dividing up the world - then the revisionist is proposing dividing up the world slightly differently. Other options are available, and we suspect you'll be able to translate what we say into the terms of your chosen view.

³⁴ In what follows we often talk of revisionary accounts "preserving" various principles. Because the relevant principles are phrased in terms of knowledge rather than knowledge*, this is just convenient shorthand. Similarly, in places we talk about "knowledge", leaving it open whether we mean the original or revised concept. In both cases the rationale is that the original and revised concepts share certain important features (more on this below).

relevant concept. But which features are we going to preserve? The two key notions are going to be principles and functions, which we'll explain in turn.

Principles are general truths that plausibly govern the relevant concept (here, knowledge). While we discussed some of the principles governing knowledge in §2, our list was incomplete. For example, we left out the principle that, if S knows p , then p is true (Factivity) and the principle that, if S knows p , then S believes p (Belief). We want to highlight that it is no part of our account that any of these principles are more fundamental, and in particular it is no part of our account that the principles are the sorts of things one has to grasp in order to be a competent user of the knowledge concept.

Turning to functions, as a rough first pass the functions of a concept X are the things that we expect X to do or bring about. In the literature one finds these functions for knowledge:

- i. to identify good or reliable informants.³⁵ Example: In saying that Catriona knows the bus leaves at 3pm I identify her as a good informant about when the bus leaves.
- ii. to indicate when an inquirer can reasonably terminate inquiry.³⁶ Example: In saying that I know the bus leaves at 3pm I signal that I can reasonably terminate my inquiry into when the bus leaves.
- iii. to provide grounds for action and sincere assertion; to provide premises for sound reasoning and deliberation.³⁷

Why think that (i)-(iii) are functions of knowledge? One argument goes via linguistic considerations. For instance, there is often something wrong about making these claims:

- (1) Catriona knows the bus leaves at 3pm, but she isn't a good informant as to when the bus leaves.
- (2) I know the bus leaves at 3pm, but I need to inquire further into when the bus leaves.
- (3) I know the bus leaves at 3pm, but I'm not going to use this as a basis for action.

Of course, there are certain contexts in which these claims are acceptable. This is perhaps most obvious with (1). Imagine that, while Catriona is well-informed about when the bus leaves, she refuses to tell anyone. One way to describe Catriona would be by uttering (1). So a nice

³⁵ See Craig, *Knowledge and the State of Nature*, Chap. 2.

³⁶ See Kappel, 'Saying That Someone Knows' Kelp, 'What's the Point'; Rysiew, 'Epistemic Scorekeeping'.

³⁷ See Bird, 'Justified Judging'; Hawthorne, *Knowledge and Lotteries*, Chap. 2, Chap. 4.

explanation of why (1)-(3) often seem wrong is that we often expect knowledge to be used to do certain things, as in (i)-(iii).

Another argument for (i)-(iii) notes that we need a concept that answers to various requirements.³⁸ First, we require a way of identifying those around us who have information about the world, given that we can't obtain all the information we need ourselves. Second, inquiry is, in principle, open-ended. No matter how much I have inquired into some matter, I can still gather more evidence. So we require a way of identifying when to terminate inquiry. Third, we require a way of identifying which beliefs we can (and which beliefs we can't) use in reasoning about how to act. There are rather stringent conditions a concept would have to meet to be a candidate for meeting these requirements. It would have to be an epistemic concept and, given the everyday nature of the requirements, it would also have to be an ordinary language term. This rules out familiar items of epistemic vocabulary such as 'justification', and leaves 'knowledge' as the obvious candidate.

These arguments notwithstanding, we don't want to commit to these being *the* functions of knowledge. It could be that there are others, and it could even be that we don't often expect knowledge to do these things (though we doubt this is the case). But we'll need to have concrete proposals for the functions of knowledge for the purposes of outlining our methodology, and (i)-(iii) are the best available candidates. So, with this proviso, we'll assume them in what follows.

Principles and functions differ in two crucial respects. First, the principles governing a concept X are things that are meant to be true of all instances of X, whereas the functions of a concept X are things we expect many (maybe all) uses of X to do or bring about. So a crucial difference between (Knowledge-Action Principle) and the third function claim is that the latter, but not the former, permits exceptions.³⁹ Second, the principles governing a concept may but need not specify its functional role. (Parity of Evidence Principle) says that two subjects alike evidentially with respect to some proposition *p* are either both in a position to know *p* or both not in a position to know *p*. This says nothing about the functional role of knowledge or knowledge ascriptions.

³⁸ See Craig, *Knowledge and the State of Nature*, Chap. 1, Chap. 2; Kelp, 'What's the Point'.

³⁹ Another is that (Knowledge-Action Principle) is a normative claim, whereas the third function claim is about the expected use of a concept.

We can now outline the constraints on our revisionary methodology. The first and most important constraint is what we call the ‘Function Constraint’:

- i. ‘Function Constraint’: The best revisionary account of knowledge* is the account that is best suited to fulfilling the functions of knowledge.

As we’ll sometimes put it, the aim is to ‘build’ our account of knowledge* around the relevant functions. Of course, it might be that there are a number of accounts of knowledge* that meet this constraint. In order to choose between them we can apply the second constraint, which we call the ‘Presumption of Innocence Constraint’:

- ii. ‘Presumption of Innocence Constraint’: If we are given no reason to reject a principle governing knowledge - such as that it is inconsistent with another principle, or prevents knowledge from performing a function - we should include it in our account of knowledge*.

The motivation for this constraint is a moderate methodological conservatism. If we have two accounts that serve the functions of knowledge equally well, we should prefer the account that is least revisionary. While this is a sort of conservatism, it only applies in a very special case, viz. when two accounts serve the functions equally well.

Putting these constraints together, the aim is to develop an account of knowledge* that is best suited to fulfilling the various functions that knowledge serves. If this doesn’t uniquely determine an account of knowledge*, we should choose the account that is least revisionary. The result of this process will be an account of knowledge* that preserves a central feature of knowledge, namely the functions that it serves. The main difference between knowledge and knowledge* will be that the latter is custom-built to serve these functions.⁴⁰

While we have been emphasising the continuities between the revisionary project and the more traditional approach, it is important to note the differences. Two are worth mentioning. First, anyone who wants to argue that knowledge is governed by a consistent sub-set of the principles

⁴⁰ While applying the mechanics of evolutionary theory to concepts is a bit of stretch, one might expect our concepts, much like our bodies, to be well adapted to our environment. So why think our knowledge concept isn’t already ideally suited to serve the relevant functions? First, as we’ll argue in the next sub-section only a subset of the principles governing knowledge are ideally suited to serve the relevant functions. Whatever forces have shaped our knowledge concept, they have added unnecessary extras. Second, if we take the evolutionary analogy seriously, we should expect our concepts to have unnecessary extras. Consider the human appendix, which we still have but serves no obvious function. Our suggestion is that certain principles governing knowledge are, in this respect, like the appendix.

needs an error-theoretic explanations. For instance, someone who wants to argue that knowledge is governed by (Parity of Evidence Principle) but not (Knowledge-Action Principle) needs an error-theoretic explanation why we mistakenly think knowledge is governed by (Knowledge-Action Principle). But someone who is engaging in the revisionary project doesn't need error-theoretic explanations. She isn't saying that knowledge is governed by some consistent sub-set of her principles. She is saying that knowledge should be governed by that consistent sub-set (or, equivalently, that knowledge* is).

Second, the revisionary project fundamentally differs from more traditional approaches with respect to its aim. Put somewhat metaphorically, our approach is a sort of *conceptual engineering*.⁴¹ In answering our question - what should knowledge be like? - we start by looking at what we expect knowledge to do for us and then ask what knowledge would have to be like in order to best do those things. Compare the engineer who wants to figure out what a bridge should do, in order to design a better bridge. She will start by looking at what good bridges are expected to do and how they do it. She can then ask what her bridge would have to be like in order to do those things better than existing bridges. In the process she might realise that her bridge should be slightly different from other bridges, because this would make it better suited to its purpose. Similarly, if our account of knowledge has to be revised in order for knowledge to best serve the relevant functions, we should make those changes.

With the methodology in place, we turn to developing a revisionary account of knowledge*.

3.2. Revisionary Sensitive Invariantism

To recap: we will construct the concept knowledge* around our three functions. If this doesn't determine a unique account of knowledge*, we'll choose the account that preserves the greatest number of principles. Our first step will be to identify two consistent subsets of the various principles. Our second step will be to identify which of these subsets is supported by our functions.

First, recall that (Social Principle) and (Knowledge-Action Principle) conflict with (Parity of Evidence Principle) and (Testimony Principle). This is not surprising. (Social Principle) relativises the rightness of ascribing knowledge to the social context, and (Knowledge-Action Principle) relativises knowledge itself to some combination of the practical environment and salience effects. In contrast, both (Parity of Evidence Principle) and (Testimony Principle) rely on the

⁴¹ We borrow this metaphor from Haslanger, 'Gender and Race'.

thought that the amount of evidence required for knowing and transmitting that knowledge remains stable and invariant through all contexts. As we'll put it in what follows, the thought is that knowledge is stable - it doesn't vary with non-evidential factors. Recall also that anything that goes for (Parity of Evidence Principle) goes for (Strength of Evidence Principle), because the former entails the latter.

Second, recall that (Odds Principle), (Weak Knowledge Principle) and (Strength of Evidence Principle) are inconsistent. The problem with this set is that it looks like we can have knowledge about certain mundane matters on fairly weak evidence, yet not have knowledge about other matters on extremely strong evidence. I have far more evidence that I will lose the lottery based on knowing the odds than I have that Manchester United won 2-1 last night based on reading the paper, yet the latter looks like knowledge whereas the former doesn't.

Finally, recall that (Odds Principle), (Weak Knowledge Principle) and (Closure Principle) are inconsistent because, via (Closure Principle), my knowing some mundane matter (e.g. I can't afford to go on safari) entails my knowing some non-mundane matter (e.g. I haven't won the lottery).

Two consistent subsets of principles are emerging. On the one hand, considerations supporting stability support (Parity of Evidence Principle), (Testimony Principle) and (Strength of Evidence Principle). This set of principles looks to be incompatible with (Knowledge-Action Principle) and (Social Principle), and one of (Weak Knowledge Principle), (Odds Principle) and (Closure Principle). We will call a view that preserves this subset of principles 'revisionary insensitive invariantism'. We use this label because it is a revisionary version of a non-revisionary view often labelled 'insensitive invariantism', on which whether a subject S knows p just depends on the usual truth-conducive factors.⁴² The main difference between this view and revisionary insensitive invariantism is that the latter is a view about knowledge*, whereas the former is a view about knowledge.

On the other hand, considerations supporting the partial relativisation of both the rightness of ascribing 'knowledge', and knowledge itself, to particular contexts support (Knowledge-Action Principle) and (Social Principle). This set of principles looks to be incompatible with (Parity of Evidence Principle), (Testimony Principle) and (Strength of Evidence Principle), and one of (Odds Principle), (Weak Knowledge Principle) and (Closure Principle). We will call a view that

⁴² See Brown, 'Contextualism and Warranted Assertibility'; Gerken, 'Warrant and Action'; Rysiew, 'Context-Sensitivity of Knowledge'.

preserves this subset of principles ‘revisionary sensitive invariantism’. We use this label because it is a revisionary version of a view often labelled ‘sensitive invariantism’, on which whether S knows p depends on a combination of truth-conducive factors, practical factors and salience effects.⁴³ Again, the main difference between this view and revisionary sensitive invariantism is that the latter is a view about knowledge*, whereas the former is a view about knowledge.

This table compares the two views:

	<i>Revisionary insensitive invariantism</i>	<i>Revisionary sensitive invariantism</i>
Accepts	<ul style="list-style-type: none"> • (Parity of Evidence Principle) • (Testimony Principle) • (Strength of Evidence Principle) • Two of: (Closure Principle), (Weak Knowledge Principle) and (Odds Principle) 	<ul style="list-style-type: none"> • (Knowledge-Action Principle) • (Social Principle) • Two of: (Odds Principle), (Weak Knowledge Principle) and (Closure Principle)
Denies	<ul style="list-style-type: none"> • (Knowledge-Action Principle) • (Social Principle) • One of: (Closure Principle), (Weak Knowledge Principle), (Odds Principle) 	<ul style="list-style-type: none"> • (Parity of Evidence Principle) • (Testimony Principle) • (Strength of Evidence Principle). • One of: (Closure Principle), (Weak Knowledge Principle), (Odds Principle)

Which revisionary view should be preferred? In what follows we argue that the functions of knowledge lead us towards revisionary sensitive invariantism, and away from revisionary insensitive invariantism.

First, two of our functions support (Knowledge-Action Principle), which says that knowing p is necessary and sufficient for treating p as a reason for action. Here are two considerations in support of the sufficiency direction. First, if knowledge has the function of providing grounds for action, the best way for knowledge to serve this function would be if knowledge were sufficient for rational action. Second, inquiry often leads to action. When I inquire into when the next bus leaves from my stop, the conclusion of my inquiry will result in action (going to the bus stop at a certain time). If knowledge has the function of indicating when an inquirer can

⁴³ See Fantl and McGrath, *Knowledge in an Uncertain World*; Hawthorne, *Knowledge and Lotteries*; Stanley, *Knowledge and Practical Interests*.

reasonably terminate inquiry, then the best way for knowledge to serve this function would be by being sufficient for rational action.

We also offer a consideration in support of the necessity direction. If knowledge were sufficient but not necessary for rational action then there would have to be some other epistemic position (or cluster of positions) with the function of providing grounds for action and terminating inquiry, at least in the cases in which knowledge is not required. But it is unclear what that position could be. Consider some candidates: believing, truly believing, justifiably believing, and truly justifiably believing. The problem is that one can (truly/justifiably/truly justifiably) believe p yet not be rational in acting on p or reasonable in terminating inquiry into whether p . I may (truly/justifiably/truly justifiably) believe my lottery ticket has lost, but that doesn't mean I am rational in throwing it away, or reasonable in not checking the results. This suggests that something stronger is needed for rational action, and the obvious candidate is knowledge.

We don't deny that knowledge *could* serve both functions while being neither necessary nor sufficient for rational action.⁴⁴ But our aim is to engineer an account of knowledge that is 'custom built' to serve a number of functions. If we are designing knowledge to serve the function of providing grounds for action and signalling when inquiry can be reasonably terminated, the best way to do so is to have an account of knowledge on which it is both sufficient and necessary for rational action and terminating inquiry. Even if there are other ways of securing the desired function, it is far better to build our account of knowledge in such a way that the function is served than to hope that features of pragmatics can do the job.

Second, the function of identifying good informants supports (Social Principle). In order to be a good informant a subject must be sensitive to the social context in which she is offering her information. Take a variant of the Jane case in §1.2. If Jane approaches Tom to ask which train to get then Tom should be sensitive to exactly how important it is that Jane get the right train. If it is imperative that Jane get on a particular train (because it's the fast train), Tom should be very careful before saying he 'knows' which train it is. In contrast, if it matters little which train Jane gets, Tom can be a little less careful before saying he 'knows' which train it is.

These two considerations support (Knowledge-Action Principle) and (Social Principle). Consequently, they support revisionary sensitive invariantism over revisionary insensitive invariantism, as the revisionary insensitive invariantist seems forced to reject both principles. So

⁴⁴ See Gerken, 'Roles of Knowledge'.

revisionary sensitive invariantism gives us an account of knowledge* that satisfies the Function Constraint. By the lights of our revisionary methodology, it is the best account of knowledge* available.

We finish this sub-section by considering two objections. First, the function of providing grounds for action supports (Weak Knowledge Principle). There are many circumstances in which we have to act on relatively weak evidence, and are clearly perfectly rational to do so. If knowledge has the function of providing grounds for action, then the fact that we are often rational in acting on relatively weak evidence suggests that we can often know on relatively weak evidence. But the revisionary sensitive invariantist has to reject one of (Weak Knowledge Principle), (Odds Principle) and (Closure Principle).

Given her revisionary aspirations, the revisionary sensitive invariantist is entitled to reject (Odds Principle) or (Closure Principle). But this might make one worry about the soundness of the methodology. So it is important to note that she has another move available. Recall the puzzle: it seems one can know one can't afford to go on safari, yet from (Closure Principle) it follows that one thereby knows one hasn't won the lottery, which contradicts (Odds Principle). The revisionary sensitive invariantist says that whether a subject knows depends on certain non-truth conducive factors, such as which error-possibilities are salient to the subject. The intuitive idea is that, the more error-possibilities are salient, the harder it is to know. As Hawthorne has argued, it is plausible that going through the process of deducing that one hasn't won the lottery changes what error-possibilities are salient in such a way that, at the end of the process, one no longer knows one can't afford to go on safari.⁴⁵ Consequently, one doesn't know one hasn't won the lottery, and (Odds Principle) and (Closure Principle) are both respected. Thus, the revisionary sensitive invariantist has the resources to deny that there is any sort of puzzle here. But she can only do so because she denies (Parity of Evidence Principle).

Second, revisionary sensitive invariantism doesn't sit well with (Testimony Principle). (Testimony Principle) says that, as long as certain conditions are met, one can gain knowledge just through being told. But knowledge can't always be transmitted across different social or practical contexts. If it doesn't matter much to Tom whether this is the fast train or not, whereas it does matter a lot to Jane, Jane can't come to know that it is the fast train just by Tom telling her that this is the fast train.

⁴⁵ Hawthorne, *Knowledge and Lotteries*, Chap. 4.

Again, the revisionary sensitive invariantist is entitled to reject (Testimony Principle). But, again, this might make one wonder about the soundness of the methodology. However, while the revisionary sensitive invariantist can't save (Testimony Principle), there is a simple modification of it that she can endorse:

(Testimony Principle*) If S knows* p , then S* would be in a position to know* p if S were to tell her so (assuming S is generally trustworthy, *the context of S* is not more demanding than that of S*, and nothing else interferes with the transmission of knowledge*).⁴⁶

(Testimony Principle*) is fully compatible with (Knowledge-Action Principle) and (Social Principle) because it restricts (Testimony Principle) to all and only those contexts that don't violate (Knowledge-Action Principle) and (Social Principle).

Here are two worries about (Testimony Principle*). First, it doesn't fit with how we use testimony. We don't generally check the context of our audiences before issuing testimony, and we don't generally check the context of the testifier before receiving it. Second, it threatens the spread of knowledge* through testimonial exchange. In all cases where the hearer is in a more demanding context than the testifier, testimony fails to transmit knowledge*.

On the first worry, (Testimony Principle*) tells us when testimonial exchanges transmit knowledge*, not when speakers will testify. Testifiers don't always know the context of their audience, and their audience doesn't always know the context of the testifier, but they testify regardless. That they do so isn't surprising. Figuring out the context of one's audience before testifying will often be impractical, difficult and not worth the effort. Much the same goes for figuring out the context of a testifier before hearing their testimony. We suspect that much of the time we operate on the assumption that someone is in a standard sort of context unless we are given reason to think otherwise. In cases where the assumption is false, testimonial exchanges may fail to transmit knowledge*.

On the second worry, (Testimony Principle*) doesn't lead to scepticism. The principle tells us that testimony won't transmit knowledge* in certain cases. Such cases include the variant of the Jane case where Tom, to whom it matters little whether this is the fast train, tells Jane, to whom it matters a lot, that this is the fast train. But, on the plausible assumption that the testifier often isn't in a less demanding situation than her audience, many cases of testimony will transmit

⁴⁶ Put very roughly, demandingness is a function of the social and practical environment, as well as the sorts of error-possibilities one is considering.

knowledge*. Imagine neither Hannah nor Sarah particularly enjoy getting wet. Hannah asks Sarah if she knows* what the weather is like outside. Sarah says “It’s raining.” Hannah asks Sarah whether she knows* it’s raining, and Sarah says “I know* because I took the rubbish out a few minutes ago.” In this case (Testimony Principle*) says that Hannah thereby comes to know* that it’s raining. Perhaps testimonial knowledge* is harder to achieve than knowledge, but it isn’t impossible. So (Testimony Principle*) avoids any sort of scepticism.

3.3. The Carnapian Argument

We have presented the core of our case for revisionary sensitive invariantism. We finish by briefly mentioning a further benefit. This benefit is distinct from the benefits discussed above, but connects with what we take to be the best motivation for revisionary projects in general, viz. our ‘engineering’ approach to philosophical analysis.

Our revisionary approach follows Carnap and Quine in treating fruitfulness as important. For Carnap, the fruitfulness of a concept is a matter of it being useful for the formulation of universal statements and/or laws.⁴⁷ One problem with this is that most concepts are fruitful in this sense. For almost any concept, there is a range of putative necessary truths about that concept, and each necessary truth furnishes universal statements and/or laws. For example, if (Odds Principle) is true then the concept knowledge is fruitful in that it is true, universally, that one can’t know one’s lottery ticket has lost just based on knowing the odds. So, unlike Carnap, we think of the fruitfulness of a concept as a matter of whether it is useful for formulating universal statements and/or laws that are themselves fruitful.⁴⁸

Because it upholds (Knowledge-Action Principle), revisionary sensitive invariantism is also fruitful in this sense. If we adopt an account of knowledge on which (Knowledge-Action Principle) is true, we are able to model the rationality of action in terms of knowledge. One way of doing this would be within a knowledge-based decision theoretic framework.⁴⁹ According to decision theory, an agent’s choice is rational when it delivers maximal expected utility. According to knowledge-based decision theory, the relevant information to be considered for calculating what choice delivers the best expected utility is provided by the total knowledge of the agent. In other words, knowledge-based decision theory requires something like (Knowledge-Action Principle). Because it endorses (Knowledge-Action Principle), revisionary sensitive invariantism

⁴⁷ Carnap, *Foundations of Probability*, 5-7.

⁴⁸ This is a partial definition of fruitfulness for concepts, not of fruitfulness itself.

⁴⁹ See Hawthorne and Stanley, ‘Knowledge and Action’, 580; Weatherson, ‘Knowledge, Bets, and Interests’.

fits nicely with knowledge-based decision theory, which in turn provides a more precise account of the intuition according to which knowledge, and only knowledge, provides grounds for action and premises for sound reasoning.

In contrast, because revisionary insensitive invariantism rejects (Knowledge-Action Principle), it isn't fruitful in this sense. If (Knowledge-Action Principle) is false, we can't model rationality in terms of knowledge, as any connections are (a) contingent and (b) complicated (see Gerken forthcoming). The revisionary insensitive invariantist may complain that (Parity of Evidence Principle), whether fruitful or not, is a true principle governing knowledge. But this complaint is beside the point, given our revisionary aims in this paper.⁵⁰

4. Conclusion

Let us take stock. We started by motivating an epistemological project - what we called a 'revisionary project' - that focuses on what knowledge should be, not what knowledge is. The motivations were twofold. First, philosophical analysis should improve our conceptual resources, not just describe them. Second, there is a *prima facie* case that our concept knowledge is inconsistent. While the case is only *prima facie*, it suffices to show that knowledge isn't quite as it should be. We then outlined how to develop a revisionary account of knowledge. The basic thought was that we make a list of the things we expect knowledge to do for us, then figure out what knowledge would have to be like in order to do those things. Finally, we argued that the upshot is what we called 'revisionary sensitive invariantism'.

At the start of this paper we noted that epistemologists have said little about what knowledge should be like. While this is partly due to a general suspicion of revisionary projects, it is also due to the lack of an established methodology for figuring out what knowledge should be like. There are lots of ways in which knowledge could be improved, and it isn't clear how to choose between

⁵⁰ Some have objected that a knowledge-based account of rationality is incompatible with Bayesian decision theory, since the latter relies on the notion of partial belief rather than knowledge (Douven, 'Knowledge and Practical Reasoning'; Schiffer, 'Interest-Relative Invariantism'). Though knowledge-based decision theory *prima facie* conflicts with credence-based decision theoretic frameworks, the two can be reconciled (Weisberg, 'Knowledge in Action'). Further, we follow Weatherson, 'Knowledge, Bets, and Interests', in thinking that, if a choice between the two decision theoretic frameworks has to be made, there are independent reasons to prefer a knowledge-based decision theory.

them. We hope we have assuaged the general suspicion of revisionary projects, at least in epistemology, and made a start on putting a methodology in place.⁵¹

References

- Bird, Alexander. 'Justified Judging'. *Philosophy and Phenomenological Research* 74, no. 1 (2007): 81–110.
- Blome-Tillmann, Michael. 'Knowledge and Implicatures'. *Synthese* 190, no. 18 (2013): 4293–4319.
- Brown, Jessica. 'Contextualism and Warranted Assertibility Manoeuvres'. *Philosophical Studies* 130, no. 3 (2006): 407–35.
- Brown, Jessica. 'Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning'. *Noûs* 42, no. 2 (2008): 167–89.
- Cappelen, Herman. 'Nonsense and Illusions of Thought'. *Philosophical Perspectives* 27, no. 1 (2013): 22–50.
- Carnap, Rudolf. *Logical Foundations of Probability*. Chicago: University of Chicago Press, 1950.
- Chalmers, David. 'Verbal Disputes'. *Philosophical Review* 120, no. 4 (2011): 515–66.
- Craig, Edward. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Oxford University Press, 1990.
- DeRose, Keith. 'Now You Know It, Now You Don't'. *The Proceedings of the Twentieth World Congress of Philosophy*, no. 5 (2000): 91–106.
- Douven, Igor. 'Knowledge and Practical Reasoning'. *Dialectica* 62, no. 1 (2008): 101–18.
- Fantl, Jeremy, and Matthew McGrath. *Knowledge in an Uncertain World*. Oxford: Oxford University Press, 2009.
- Fisher, Justin. 'Meanings and Methodologies', in *New Waves in Philosophy of Mind*, ed. Jesper Kallestrup and Mark Sprevak (London: Palgrave MacMillan, 2014), 54–76.
- Gerken, Mikkel. 'The Roles of Knowledge Ascriptions in Epistemic Assessment'. *European Journal of Philosophy* 23, no. 1 (2015): 141–61.
- Gerken, Mikkel. 'Warrant and Action'. *Synthese* 178, no. 3 (2011): 529–47.
- Haslanger, Sally. 'Gender and Race: (What) Are They? (What) Do We Want Them to Be?'. *Noûs* 34, no. 1 (2000): 31–55.

⁵¹ We would like to thank Michael Hannon, Nick Hughes, Sebastian Köhler, Martin Kusch, Jacques Vollet, Stephen Wright, and audiences at the University of Edinburgh and the University of Hradec Králové. Davide Fassio's work on this paper was supported by the SNSF research project "Knowledge-based Accounts of Rationality". Robin McKenna's work was supported by funding from the ERC Advanced Grant Project "The Emergence of Relativism" (Grant No. 339382) and a postdoctoral fellowship from the University of Geneva.

- Haslanger, Sally. 'What Good Are Our Intuitions? Philosophical Analysis and Social Kinds'. *Proceedings of the Aristotelian Society* 106 (2006): 89–118.
- Haslanger, Sally. 'What Knowledge Is and What it Ought to Be: Feminist Values and Normative Epistemology'. *Philosophical Perspectives*, 13 (1999): 459–80.
- Hawthorne, John. *Knowledge and Lotteries*. Oxford: Oxford University Press, 2004.
- Hawthorne, John, and Jason Stanley. 'Knowledge and Action'. *Journal of Philosophy*, 105 no. 10 (2008): 571–90.
- Jones, Ward. 'Why Do We Value Knowledge?'. *American Philosophical Quarterly* 34, no. 4 (1997): 423-39.
- Kappel, Klemens. 'On Saying That Someone Knows: Themes from Craig', in *Social Epistemology*, eds. Adrian Haddock, Alan Millar, and Duncan Pritchard (Oxford: Oxford University Press, 2010): 69-88.
- Kelp, Christoph. 'What's the Point Of "Knowledge" Anyway?'. *Episteme* 8, no. 1 (2011): 53–66.
- Kusch, Martin. *Knowledge by Agreement: The Programme of Communitarian Epistemology*. Oxford: Oxford University Press, 2002.
- Littlejohn, Clayton. 'The Unity of Reason', in *Epistemic Norms: New Essays on Action, Belief, and Assertion*, eds. Clayton Littlejohn and John Turri (Oxford: Oxford University Press, 2014): 135-54.
- Pritchard, Duncan. 'On Meta-Epistemology'. *The Harvard Review of Philosophy* 18, no. 1 (2012): 91-108.
- Quine, Willard van Orman. *Word and Object*. Cambridge, MA: MIT Press, 1960.
- Riggs, Wayne. 'The Value Turn in Epistemology', in *New Waves in Epistemology*, ed. Vincent Hendricks (London: Palgrave Macmillan, 2008): 300-23.
- Rysiew, Patrick. 'The Context-Sensitivity of Knowledge Attributions'. *Noûs* 35, no. 4 (2001): 477–514.
- Rysiew, Patrick. 'Epistemic Scorekeeping', in *Knowledge Ascriptions*, eds. Jessica Brown and Mikkel Gerken (Oxford: Oxford University Press, 2012): 270-93.
- Schiffer, Stephen. 'Contextualist Solutions to Scepticism'. *Proceedings of the Aristotelian Society* 96 (1996): 317-33.
- Schiffer, Stephen. 'Interest-Relative Invariantism'. *Philosophy and Phenomenological Research* 75, no. 1 (2007): 188–95.
- Smith, Martin. 'What Else Justification Could Be'. *Noûs* 44, no. 1 (2010): 10-31.
- Stanley, Jason. *Knowledge and Practical Interests*. Oxford: Oxford University Press, 2005.

- Vogel, Jonathan. 'Are There Counterexamples to the Closure Principle?', in *Doubting: Contemporary Perspectives on Skepticism*, eds. Michael D. Roth and Glenn Ross (Dordrecht: Kluwer, 1990): 13-29.
- Weatherson, Brian. 'Knowledge, Bets, and Interests', in *Knowledge Ascriptions*, eds. Jessica Brown and Mikkell Gerken (Oxford: Oxford University Press, 2012): 75-103.
- Weiner, Matt. The (Mostly Harmless) Inconsistency of Knowledge Ascriptions. *Philosophers' Imprint* 9, no. 1 (2009): 1-25.
- Weisberg, Jonathan. 'Knowledge in Action'. *Philosophers' Imprint* 13, no. 22 (2013): 1-23.
- Williamson, Timothy. *Knowledge and its Limits*. Oxford: Oxford University Press, 2000.
- Yourgrau, Palle. 'Knowledge and Relevant Alternatives'. *Synthese* 55, no. 2 (1983): 175-90.