Misleading Evidence and the Dogmatism Puzzle

Ru Ye

Australasian
Association of
Philosophy

Routledge
Taylor & Francis Group

Misleading Evidence and the Dogmatism Puzzle

Ru Ye

Cornell University

ABSTRACT
According to the Dogmatism Puzzle presented by Gilbert Harman, knowledge induces dogmatism because, if one knows that p, one knows that any evidence against p is misleading and therefore one can ignore it when gaining the evidence in the future. I try to offer a new solution to the puzzle by explaining why the principle is false that evidence known to be misleading can be ignored. I argue that knowing that some evidence is misleading doesn't always damage the credential of the evidence, and therefore it doesn't always entitle one to ignore it. I also explain in what kind of cases and to what degree such knowledge allows one to ignore evidence. Hopefully, through the discussion, we can not only understand better where the dogmatism puzzle goes wrong, but also understand better in what sense rational believers should rely on their evidence and when they can ignore it.

1. Introduction

It is generally claimed that we should follow evidence in belief or credence revision. We should believe what evidence supports and we should proportion credence with the degree of support. When we receive new evidence in considering whether a proposition is true or how likely it is to be true, we should consider how the evidence tips the evidential scale and we should adjust beliefs or credence accordingly. And the motivation behind this advice—'Follow your evidence'—is intuitive: We follow evidence because evidence tends to lead us to the truth.

However, this rationale also underlies our intuition that the general advice admits exception. For example, if you know that the new evidence you just acquired is misleading, it seems that you should just ignore it. We follow evidence because we want to get to the truth or at least get closer to it. If so, it seems irrational to take into account evidence that you know leads you to falsehood. Doing so would be like embarking on a road that you know leads you to the north when you want to go to the south. This intuition can be framed into the following principle [Lasonen-Aarnio 2014: 419]:

> *Entitlement*
> For any proposition *p*, time *t*, person S, and piece of evidence *e*, if S knows at *t* that *e* is misleading with regard to *p* then at *t* S could ignore *e* as it bears on *p*.

Entitlement sounds plausible. You could rely on evidence that you have no idea is misleading. But it seems silly to rely on evidence that you *know* is misleading. As Lasonen-

Aarnio [ibid.: 420] says, 'Taking into account evidence one knows to be misleading seems like relying on the testimony of a person one knows to be a pathological liar.'

As intuitive as *Entitlement* sounds, it leads to an unacceptable dogmatist conclusion that one's knowledge that *p* enables one to ignore all future evidence against *p*. The puzzle originates from Saul Kripke and is sharply presented by Gilbert Harman as follows [1973: 148−9]:

> If I know that h is true, I know that any evidence against h is evidence against something that is true; so I know that such evidence is misleading. But I should disregard evidence that I know is misleading. So, once I know that h is true, I am in a position to disregard any future evidence that seems to tell against h.

Harman calls this 'the dogmatism puzzle': one's knowledge that h seems to enable one to dogmatically ignore all new evidence against h.

This paper offers a solution to the puzzle by explaining why *Entitlement* is false. In Section 2, I present more precisely the dogmatism puzzle. In Section 3, I argue that the puzzle can be solved only in a defeasibilist framework. In Section 4, I explain why *Entitlement* is false in that framework. I argue that knowing that some evidence is misleading doesn't always enable one to ignore it, because such knowledge doesn't always damage the credential of the evidence. Section 5 deals with the worry that denying *Entitlement* allows knowingly moving further away from the truth. Finally, in Section 6, I explain in what kind of cases knowing that some evidence is misleading entitles one to ignore it. Hopefully, through the discussion, we can not only understand better where the dogmatism puzzle goes wrong, but also understand better in what sense rational believers should rely on their evidence and when they can ignore evidence.

2. The Dogmatism Puzzle

The dogmatism puzzle can be presented as follows. Suppose that

(1) at a time t_0, a person S knows that *p* (for some proposition *p*).

Suppose that S infers '*p* is true' from *p*. Since 'misleading evidence' just means evidence that supports a falsehood or that counts against a truth, we can suppose that from '*p* is true' S infers that any evidence against *p* is misleading. So,

(2) at t_0, S knows that any evidence against *p* is misleading.[1]

---

[1] Some scholars think that the dogmatist's reasoning is already faulty at this step. They see this as a case of closure failure (see Cargile [1995]; Sharon and Spectre [2010]). I find this diagnosis unconvincing, because typical factors that threaten closure don't need to play a role here. For example, S's inference from 'p' to 'any evidence against p is misleading' doesn't need to involve risk-accumulation through premises-conjunction: We can stipulate that S is rationally certain that 'any evidence against a truth is misleading', given that this is true by the definition of 'misleading evidence'. S's inference doesn't need to involve context shift—we can stipulate that S is thinking about *p* and about evidence against *p* in a fixed context. It also doesn't involve inference from a falsehood. Sharon and Spectre [ibid.] think that S's inference is based on the false belief that *e* won't obtain at t_1. But S's belief is not (or doesn't need to be) based on this false belief at all; it is based just on her belief that *p* and her belief that any evidence against a truth is misleading. With that said, when *e* is strong evidence against *p*, it does sound a little awkward for S to assert at t_0 'If *e* is evidence against *p*, then *e* must be misleading.' However, as Sorensen [1988] and Hawthorne [2004: 71−3] argue, the awkwardness is explained by the fact that S's knowledge of the conditional at t_0 is junk: if S comes to know the antecedent at t_1, S would no longer know the conditional and therefore would not know its consequent.

Suppose that, at a later time t_1, S acquires evidence e, which she knows is evidence counting against p. Then

(3) at t_1, S knows that e is misleading.

Given *Entitlement*, it follows that

(4) at t_1, S can ignore e.

Discharging all suppositions, we seem to have the following conclusion:

> *Knowledge-Dogmatism*
> If S knows that p at a time (and if S is capable enough to know that if p then any evidence against p is misleading; this complication is omitted hereafter), then, when she later acquires evidence that she knows counts against p, she could just ignore it.

Here is an example drawn from Harman [ibid.: 148]. Suppose I know that Tom stole the book from the library. I then know that any potential evidence against 'Tom stole the book from the library' is misleading. For example, I know that if Tom's mother (whom I know is very reliable) were to tell me that it was Tom's twin brother who stole the book, this would be misleading evidence. At a later time, Tom's mother tells me that it was Tom's twin brother who stole the book. But since I know that this testimony is misleading, I can just ignore it.

The dogmatist's argument looks cogent. And yet the conclusion *Knowledge-Dogmatism* is problematic because, intuitively, knowing p shouldn't put you in a position to ignore all future evidence against p. It seems that, even if you know p now, when you obtain evidence against p in the future you should still take it into account, and you should be prepared to correct your opinion if that evidence is strong. In the case described above, it seems that I should not just ignore it when Tom's mother tells me that it was Tom's twin brother who stole the book. Instead, I should adjust my belief or credence according to how this testimony affects my total evidential situation.

3. Why Denying Indefeasibilism Is Necessary to Solving the Puzzle

Harman's solution to the puzzle is to identify and deny the indefeasibilism assumption in the dogmatist's reasoning. (I will clarify the term 'indefeasibilism' later on; for now, we can just understand it as the view that one's knowledge that p is indefeasible by any new evidence against p.) Note that there is a gap from (2) to (3) in the dogmatist's reasoning. That S knows at t_0 that any evidence against p is misleading doesn't imply that S knows at t_1 that e is misleading. The gap is bridged only if S *still* knows at t_1 that any evidence against p is misleading. This condition is met only if S still knows p at t_1. After all, S has no independent evidence that e is misleading. But to say that his knowledge that p remains despite any new evidence against p is to assume indefeasibilism. So, if we deny indeafeasibilism, the dogmatist reasoning falls apart.

However, as Lasonen-Aarnio [2014] correctly criticizes, simply denying indefeasibilism is not enough to solve the puzzle. Even if knowledge is sometimes defeated by new evidence, it's not always so: there are cases in which the new evidence is not strong enough to defeat S's knowledge. In those cases, S still knows that the new evidence e is misleading and thus (by *Entitlement*) S can ignore e. This means that S doesn't need to reduce his credence in p.[2] But this is still counterintuitive; for even though e is not

strong enough to defeat S's knowledge p, this doesn't mean that S's credence can be entirely unmoved by e. This counterintuitive conclusion cannot be avoided just by denying indefeasibilism. To say that S cannot ignore e because it might defeat knowledge cannot explain why S cannot ignore e when the knowledge is not defeated.

Now, although I agree with Lasonen-Aarnio that denying indefeasibilism is not enough to solve the puzzle, I do think it is necessary. To see this, it is important to be clear on exactly what kind of indefeasibilism is assumed by the dogmatist. Compare the following two versions of indefeasibilism:

> *Extreme Indefeasibilism*
> Once one knows that p, any further evidence one gains against p cannot defeat this knowledge.
> *Moderate Indefeasibilism*
> Once one knows that p, any further evidence one gains against p cannot defeat this knowledge *simply by reducing the evidential probability of p below the threshold required for knowledge.*

The dogmatist must not assume *Extreme Indefeasibilism*, because it is just too implausible. Ram Neta [2009: 170, 182], a prominent indefeasibilist, argues that, even if one knows p, one could lose the belief p, or the old evidence one has for p, or the ability to base belief on the old evidence, as a result of gaining new evidence against p. Consider this example. You once knew that your brain tumour is malignant, based on the testimony of a prestigious oncologist. But, immediately later, your family doctor tells you that the tumour is benign. You are so thrilled by this testimony that you totally forget the earlier one given by the oncologist. So, you no longer know that the tumour is malignant. But here, you lose knowledge merely because you lose your old evidence as a result of gaining the new piece, not because the evidential power of the old evidence is overridden by that of the new piece.

Since this kind of examples is perfectly possible, even an indefeasibilist should allow knowledge to be defeated by new evidence (although he might call this, not 'defeat', but instead 'losing knowledge as a result of gaining new evidence'). What an indefeasibilist shouldn't allow, according to Neta, is that one's knowledge is defeated by new evidence *simply because* the evidence 'tips the evidential scale' so that the evidential probability is below the threshold required for knowledge. So, in the above example, if you still remember the old testimony (and if you maintain your belief on that basis), the indefeasibilist would say that your knowledge is not defeated by the new evidence: if the evidential force of the old evidence is indeed strong enough for you to have knowledge earlier, then gaining new evidence cannot undermine that force to the extent that you lose the knowledge.

So, the dogmatist must assume something weaker than *Extreme Indefeasibilism*—something like *Moderate Indefeasibilism* instead. But if so, the conclusion to draw from his reasoning should be something weaker than *Knowledge-Dogmatism*. If the dogmatist admits that one can lose knowledge as a result of gaining new evidence, his conclusion should be weakened to this:

---

2 As Lasonen-Aarnio [2014: 423−4] points out, the term 'ignore' in *Entitlement* should be understood as 'completely ignore': one ignores e as it bears on p when one takes e into account *neither* in the revision of one's belief toward p nor in the revision of one's credence in p. A reason for this reading is that *Entitlement* would be unmotivated if it merely says that knowing that e is misleading enables one to ignore it in belief revision but is silent on whether it enables one to ignore e in credence revision. Given the close connection between credence and belief, credence revision could sometimes bring about or demand belief revision. Therefore, it is hard to maintain that knowing that e is misleading always enables one to ignore it in belief revision but doesn't enable one to ignore it in credence revision.

Knowledge-Dogmatism[*]

Once one knows that p, one can ignore any new evidence that one knows counts against p, *if the new evidence doesn't result in one's loss of the relevant belief or old evidence or basing ability.*

Although this conclusion is weaker than *Knowledge-Dogmatism*, it is still unacceptable by our intuition; therefore, the dogmatism reasoning still presents us with a puzzle. Consider, again, Harman's stolen-book example. Knowing that Tom stole the book doesn't seem to entitle me to ignore the testimony from Tom's mother, even if we stipulate that acquiring the testimony doesn't result in the loss of my belief or old evidence or ability to base belief on that evidence.

Now, we can see why denying indefeasibilism is necessary to solve the puzzle (hereafter, when I mention 'indefeasibilism' I mean *Moderate Indefeasibilism*, since *Extreme Indefeasibilism* is just too implausible). This is because an indefeasibilist is committed to *Knowledge-Dogmatism*[*] even without the dogmatist's reasoning—especially without detouring through *Entitlement*. Let me explain.

As mentioned above, an indefeasibilist denies that new evidence against p can defeat one's knowledge p by dragging the evidential probability of p below the threshold for knowledge. But an indefeasibilist must hold that knowledge requires probability 1.[3] If so, the indefeasibilist must maintain that if one knows that p, and if one's new evidence against p doesn't result in a loss of one's old evidence, then the new evidence cannot reduce the evidential probability of p—not even a little: If $\Pr(p/E) = 1$ (assuming that E is one's total evidence), then when one gains new evidence e and this doesn't result in loss of the old evidence E, the evidential probability of p would be given by $\Pr(p/E\&e)$, which would still be 1, according to probability axioms. Then *Knowledge-Dogmatism*[*] follows. If gaining new evidence e doesn't result in one's loss of the relevant belief or old evidence or basing ability, and if it doesn't make p any less likely, why should you revise belief or credence in light of e?

So, even without the dogmatist's argument from *Entitlement*, *Knowledge-Dogmatism*[*] is unavoidable if indefeasibilism is true. This means that denying indefeasibilism is necessary to solving the puzzle. So, in my later discussion of why *Entitlement* is false and in my answer to objections, I will be working under a defeasibilist picture of knowledge. That is, I will assume that new evidence can sometimes defeat knowledge by dragging the evidential probability below the threshold required for knowledge. And

---

[3] In fact, we can show— as follows— that indefeasibilism (again, by this I mean *Moderate Indefeasibilism*) is equivalent to the probability-1 requirement. (1) The probability-1 requirement implies indefeasibilism. Suppose that one knows that p based on one's evidence E. If $\Pr(p/E) = 1$, then, when one gains new evidence e and this doesn't result in the loss of one's old evidence E, the evidential probability of p is $\Pr(p/E\&e)$, which would still be 1. So, one's knowledge will not be defeated by e. (2) If the probability-1 requirement is false, indefeasibilism would be false. Here is a proof adapted from Neta [2009: 175−6]. Suppose that one knows that p based on one's evidence E. If knowledge doesn't require probability 1, we can further suppose that $\Pr(p/E) \neq 1$. Then $\Pr(\text{not-}p/E) > 0$. If so, there must be a proposition Z that is true and that is extremely unlikely, given E, such that $\Pr(Z/E)$ is far smaller than $\Pr(\text{not-}p/E)$. (In Neta's example, Z is a true proposition specifying precisely the distribution of microphysical states over all points in space−time.) Now consider a proposition e, which is 'Z or not-p': e must be true, because Z is true. Suppose that one acquires e as new evidence and that this doesn't result in the loss of one's old evidence E. Then the new evidential probability of p is given by $\Pr(p/E\&e)$. But notice this:

$$\begin{aligned} \Pr(p/E\&e) &= \Pr(p/E\&(Z \text{ or not-}p)) = \Pr(p\&(Z \text{ or not-}p)/E)/\Pr(Z \text{ or not-}p/E) \\ &= \Pr(p\&Z/E)/\Pr(Z \text{ or not-}p/E) \\ &\leq \Pr(Z/E)/\Pr(\text{not-}p/E) \end{aligned}$$

This ratio will be very small, because—by supposition—$\Pr(Z/E)$ is far smaller than $\Pr(\text{not-}p/E)$. So, $\Pr(p/E\&e)$ will not be high enough for knowledge, which means that e defeats one's knowledge that p.

this would amount to the assumption that knowledge doesn't require probability 1 (see note 3).

4. Why *Entitlement* Is False

In this section, I explain why *Entitlement* is false, by arguing for the following two claims. (1) In many cases of belief or credence revision, we rely on evidence not because we take what the evidence supports to be true, but because the evidence makes the proposition in question more (or less) likely to be true. (2) When we know that some evidence is misleading, this knowledge doesn't always damage the credential of the evidence—the evidence can still make the proposition more (or less) likely to be true. Therefore, we should still rely on the evidence even when we know that it is misleading.

4.1 Relying on Evidence in Probability Assessment

In some cases of belief revision, relying on evidence means believing what the evidence supports, and this requires taking what the evidence supports to be true. For example, I believe that my friend John is sick, because he just told me so. In relying on his testimony, I take what he told me to be true. Another example: I believe that it is raining outside, because I am looking outside now and I have the visual experience of its raining. In relying on the experience, I take what it supports to be true.

In these cases, the evidence in question is the total evidence that I have about the proposition, and I rely on the total evidence by believing what it supports. In so far as I aim to form a true belief in these cases, I must take what the evidence supports to be true in order to rely on it.

However, when evidence e is not our total evidence, or when our concern is not belief revision but instead credence revision, we don't need to take what e supports to be true in order to rely on it. In these cases, we rely on e in assessing the probability of the proposition in question, and we adjust belief or credence in light of the assessment. In our decision of whether to rely on e, what matters is not whether e supports a truth but, rather, whether it affects the probability of the proposition in question.[4]

Consider this example of rational belief revision. Suppose that at time t your total evidence about p consists of two independent pieces of evidence, e_1 and e_2. After carefully examining the two pieces of evidence, you find that e_1 supports p, e_2 supports not-p, and that e_1 supports p roughly to the degree to which e_2 supports not-p. So, you decide to suspend judgment on p.

---

[4] It might seem trivially true that evidence, even when it is known to be misleading, affects the probability of the proposition in question: if e is evidence about p, then surely e affects the probability of p—after all, this is a popular definition of evidence. But that thought is misguided. e is evidence about p only in the sense that it affects the probability of p in the *prior probability function* (Pr_0) (i.e. $Pr_0(p/e) \neq Pr_0(p)$); and when we ask whether knowledge that e is misleading allows you to ignore e in probability assessment, the probability function that we use is your *evidential probability function* (Pr)—namely, the probability function with your current evidence built into it. The fact that e affects the probability of p in Pr_0 doesn't automatically mean that it does so in Pr. For example, if gaining evidence e for p makes you lose some evidence for p, then gaining e might not strengthen the overall case for p and therefore might not make p more probable in Pr. Another possibility is that Pr contains information that damages the credential of e, information that Pr_0 doesn't have. For example, the fact that a reliable journal like *Times* reports that p is evidence for p in Pr_0 but, if your current evidence includes the information that *Times*'s report is based on an unreliable source this time, it might not increase the probability of p in Pr.

In reaching the rational decision to suspend judgment on p, you have relied on both e_1 and e_2. But you rely on e_1 not because you take what it supports to be true—you know that it supports p and yet you don't think p is true. The same thing can be said for your reliance on e_2. Instead, you rely on e_1 and e_2 because *they affect how likely p is true*. Since you find that e_1 supports p roughly to the same degree to which e_2 supports not-p, you conclude that p is not more (or less) likely to be true than not-p, and you decide to suspend judgment on p. So, you rely on e_1 (and also e_2) because it affects the probability of p, and probability assessment constrains your belief revision in this case. Evidence is worthy not just because it tells us whether the proposition in question is true, but also because it is the ground on which we assess how likely it is that the proposition is true.

This fact is more salient when it comes to credence revision. In deciding how to revise credence when receiving new evidence, our concern is not with whether what the evidence supports is true; rather, it is with the question of how the evidence affects the probability of the proposition in question. Even if we know before the revision that responding to the new evidence wouldn't give us a true belief (because we might know before the revision that we wouldn't form a belief at all), we still take the evidence into account in credence revision. For example, when I receive some insubstantial evidence e supporting p—insubstantial, in the sense that it won't enable me to form a belief-attitude toward p—I still take it into account in deciding how to revise credence in p. I don't just push it aside, remaining unmoved in my credence revision, when I realize that responding to e wouldn't give me a true belief.

4.2 Knowledge that e Is Misleading Doesn't Damage Its Credential

We have seen that, in credence revision and in some cases of belief revision, we rely on evidence not because we take what it supports to be true, but because the evidence affects the probability of the proposition in question, and probability assessment constrains the revision of our cognitive attitude. Given this role of evidence in probability assessment, we can see why evidence that is known to be misleading cannot always be ignored. This is because knowing that e is misleading doesn't always damage its credential. In a defeasibilist framework, e could still affect the probability of the proposition in question, even if e is known to be misleading.

Often, evidence e can make a proposition p more (or less) likely to be true only given certain background conditions. For example, your friend's testimony that he is sick makes it more likely that he is sick, only given the background condition that he is reliable. At least, your background knowledge shouldn't include information that undermines the credential of the evidence. For example, if you have visual experience of a red table but your background knowledge includes the information that the table is illuminated by red light, then your visual experience doesn't make it more likely that the table is red.

Now, our intuition for *Entitlement* might come from treating knowledge that e is misleading as relevant background information about the 'quality' or 'credential' of e. The intuition is that knowing that e is misleading damages the credential of e, just like knowing that your friend is drunk damages the credential of his testimony. And just like you should not rely on his testimony when you know that this friend is drunk, you should not rely on e when you know that it is misleading.

This intuition is mistaken. Knowing that *e* is misleading differs from knowing that your friend is drunk. When you know that your friend is drunk when he tells you 'I am sick', the knowledge is indeed relevant background information that damages the credential of the testimony, because it means that his testimony doesn't make it more (or less) likely that he is sick. (Later on, I will say more about to what degree the background information damages the credential.) However, when you know that *e* is misleading with respect to *p*, the credential of *e* might not be damaged in any way, because *e* could still affect the probability of *p*. Consider this example.

SWIM
You know that your friend Philip would go occasionally to the gym for swimming. Today you just saw that (*e*) he went into the gym for swimming. You are considering the proposition that (*p*) Philip is going to be drowned from swimming. You know that this proposition is false, because you know that Philip is a swimming champion. Given his excellent swimming skills, he would not be drowned. You conclude that *e* is misleading—it supports *p* (although only slightly), but *p* is false.

In SWIM, you know that *e* is misleading, because you know that Philip wouldn't be drowned. However, learning that Philip went swimming still raises the probability, although only very slightly, that he would be drowned when swimming. If he hadn't gone swimming, then certainly he wouldn't be drowned when swimming. But if he went swimming, there is at least a slight chance that he would be drowned. More importantly, these facts remain even if you know that he wouldn't be drowned—going swimming raises the probability of drowning, even for a swimming champion.[5] Therefore, your knowing that *e* is misleading doesn't damage the credential of *e* when it comes to affecting the probability of *p*.

In SWIM, *e* is only insubstantial evidence for *p*. But the point remains even if we make *e* more substantial. For example, you know not only that Philip went swimming but also that the tiles in the pool are especially slippery today and that the water in the pool is contaminated by some chemicals that cause people to feel dizzy. Now, *e* is substantial evidence for Philip's drowning. Had you not known that he is a swimming champion, it would be rational for you to believe that Philip is going to be drowned. However, the point remains that even though you know that he would not be drowned and therefore you know that *e* is misleading, this knowledge doesn't damage the credential of *e*: *e* still raises the probability of Philip's drowning. If you have *e* beforehand and you see Philip going to the gym, you probably would give him a friendly reminder and say, 'Be careful today, the pool is slippery and pool water is contaminated; not that you would be drowned—I know you are an excellent swimmer, but I just want to give you this note of caution.' If your knowledge that *e* is misleading damages the credential of *e*, so that *e* doesn't make *p* more probable, this reminder would be completely redundant. But this reminder makes perfect sense, and Philip would probably appreciate it.

Here are more examples to explain the above point—that knowing that *e* is misleading with respect to *p* is compatible with *e* still affecting the overall probability of *p*. Flipping a coin three times that you know is fair and learning that it lands heads each time might not deprive you of knowledge that the coin is fair. Therefore, you still know that

---

[5] Here, I assume that my knowledge that Philip wouldn't be drowned doesn't come with probability 1. Otherwise, learning that he went swimming wouldn't increase the probability of drowning. This assumption is legitimate. As I explained in section 3, I am working under a defeasibilist framework, in which knowledge doesn't require probability 1. (See note 3.)

the coin is fair and that the landing result is misleading. However, this knowledge doesn't damage the credential of the landing result, for it still makes it more likely that the coin is not fair. Learning that only 1% of the people in the city drive a Honda might not deprive you of knowledge that the suspect of the criminal case drove a Honda (because you have strong evidence, gained from witness reports and car experts' verdicts, that the suspect drove a Honda); but it still decreases the probability that the suspect drove a Honda.

These examples show that knowing that e is misleading is not some higher-order background information about the quality or credential of e. Knowing that e supports a falsehood (or counts against a truth) with respect to p doesn't undermine its role in affecting the probability of p. If so, the knowledge doesn't entitle you to ignore e—as I have argued in section 4.1, we sometimes rely on evidence exactly because of the role it plays in affecting probabilities. Therefore, having knowledge that e is misleading doesn't mean that there is anything wrong with taking e into account. After all, it is not as if the evidence only apparently affects the probability of p while in fact it doesn't. If it affects the probability of p, you are not 'deceived' or 'misled' in any way if you take it into account in assessing the probability of p and if you rely on that assessment in your belief or credence revision.[6]

5. A Worry: Moving Further Away from the Truth?

To sum up, this is my explanation of why *Entitlement* is false. Sometimes, relying on evidence doesn't require taking to be true what it is that the evidence supports. When e is not your total evidence, or when your concern is how to revise credence, relying on e doesn't require you to take to be true what e supports; it requires just that e affect the probability of p. Since e could still affect the probability of p even when you know that it is misleading, the knowledge doesn't always entitle you to ignore it. Therefore, *Entitlement* is false.[7]

However, there is a worry that the conclusion is derived too quickly. That worry has to do with the motivation for *Entitlement* mentioned at the beginning of this paper. It is true that, sometimes, relying on e doesn't require taking to be true what e supports. However, it seems that, in relying on e, *you shouldn't take e to be leading you further away from the truth*. But relying on evidence known to be misleading is knowingly moving further away from the truth.

This worry can be expressed more precisely when it comes to credence revision, where it is clear what 'moving further away from the truth' means. It is widely accepted

---

[6] Peter Baumann [2013: 14−15] argues that knowing that e is misleading doesn't entitle one to ignore e because, even if one knows that e is misleading, one might not know that one knows, or one might not know that one knows that one knows, etc. I am not convinced by this diagnosis of *Entitlement* because, even if one has infinitely higher level knowledge that e is misleading, one might still not ignore e. For there might be cases in which even the infinitely higher level knowledge doesn't damage the credential of e. Besides, to say that what's at issue is the lack of infinitely higher level knowledge implies that finite human beings can never ignore evidence; but this is counterintuitive when I have knowledge that completely damages the credential of the new evidence. For instance, knowing with probability 1 that my friend is lying seems to allow me to ignore his testimony, even if I don't have infinitely higher level knowledge.

[7] This diagnosis of *Entitlement* also clarifies in what sense knowledge allows us to end inquiry. Knowledge allows us to end inquiry only in the sense that it allows us to stop looking for further evidence, not in the sense that it allows us to ignore the evidence we just acquired. Evidence we don't have yet has no bearing on evidential probabilities, but evidence we just acquired does (for evidential probability just means probability given the total evidence we already have). Much thanks to a referee for bringing this issue to my attention.

that the goal of credence is accuracy, just like the goal of belief is truth. The basic idea of accuracy is connected to truth: If p is true, then the higher your credence in p is, the more accurate that credence is; if p is false, then the lower your credence in p is, the more accurate your credence is. So, in credence revision, moving further away from the truth means adopting a less accurate credence.

Given this accuracy goal, the worry can be expressed as follows. If you know that e is misleading with regard to p but you still revise your credence in the light of e, you will knowingly make your credence in p less accurate. Suppose that you know that e is misleading, because you know 'e counts against p, but p is true.' Suppose that you find that e makes p a little less probable and then you reduce credence in p. Then you would know that you are moving further away from the truth—knowing that 'p is true' and yet reducing credence in p is knowingly making your credence in p less accurate. But it seems irrational in credence revision to move knowingly to a less accurate credence, just like it is irrational in belief revision to move knowingly to a false belief. Therefore, it seems irrational to revise your credence in the light of e when you know that e is misleading.

Now, on the one hand, the intuition seems correct that you shouldn't knowingly move to a less accurate credence. On the other hand, the intuition is deeply problematic. First, it implies that sometimes our credence revision is not constrained by assessment of probability. It implies that even if you correctly assess that e decreases the probability of p, then—as long as the decrease is not substantial enough to deprive you of knowledge—you should not reduce credence in p. But it seems irrational to retain credence while admitting that the new evidence e makes p less likely.

Second, the intuition implies that, when you know that p and infer that p is true, you should become certain of p. If you know that p is true, you would know that certainty in p is the most accurate credence; then it seems irrational to remain, say, 0.9 confident in p. But the claim that knowledge implies certainty induces scepticism. Given that we are rarely certain of anything, the claim would imply that we rarely know anything.[8]

So, there must be something wrong with the intuition that we shouldn't knowingly move to a less accurate credence. Perhaps this intuition is based on an intuitive norm of decision-making:

K-norm
If you know that p, and if you know that if p then doing B would have less utility than doing A and that no other option has greater utility than doing A, then you should decide to do A and not to do B.

K-norm explains our intuition that we shouldn't knowingly move to a less accurate credence. In credence revision, utility is measured by accuracy, so your goal of maximising utility is that of maximising accuracy. *K-norm* implies that, if you know that p and if you know that reducing credence in p when p is true would result in a less accurate credence and that certainty is the most accurate, then you shouldn't reduce credence and you should instead boost credence in p to certainty.

---

[8] Note that even Timothy Williamson [2000: 213−14], who thinks that knowledge requires probability 1, denies that knowledge implies certainty. He denies the common association between probability 1 and certainty.

Although *K-norm* seems plausible, it conflicts with the standard norm in decision-making, the norm of maximising expected utility (*MEU*). For some actions *A* and *B*, even if you know that doing *A* has the greatest utility, doing *A* might have smaller expected utility than doing *B*. In these cases, *K-norm* recommends doing A, while *MEU* recommends not doing A. And it seems that, in cases of conflict, the recommendation given by *MEU* better matches our intuition.

Consider this example. You are offered a bet on the proposition that you have hands: You gain two dollars if it is true and you die otherwise. You know that you have hands, so you know that accepting the bet has the greatest utility (you get two dollars by accepting the bet and nothing if you reject it). *K-norm* recommends accepting the bet. But I trust that our intuition is clear that you shouldn't do so. This intuition is explained by *MEU*: when the probability of your having hands is not 1, accepting the bet means risking your life for just two dollars; therefore, it would have much smaller expected utility than rejecting the bet.[9]

Now, I don't mean to claim that this example is a knockdown argument against *K-norm*. Perhaps *K-norm* has virtues that *MEU* doesn't. Or perhaps there is no single norm that applies in every situation. What I maintain is just that we should not be so worried by the fact that respecting evidence known to be misleading implies knowingly moving away from the truth. As the above example shows, sometimes when we have knowledge but our evidence doesn't rule out all error possibilities, we can knowingly move away from the goal of maximising utility in order to strike the best balance between expected gain and loss. And as various formal epistemologists have argued, matching credence with probability is maximising expected accuracy and therefore is striking the best balance between expected gain and loss in credence revision.[10]

6. Induced Credential-Damaging Knowledge

I have argued that knowing that *e* is misleading doesn't always damage its credential and therefore doesn't always entitle you to ignore it. However, this is not to say that such knowledge *never* entitles you to ignore new evidence. In this section, I explain what the exceptions are.

Sometimes, we can gain knowledge about *e* that *does* damage its credential (in what follows, I call it 'credential-damaging knowledge' about *e*). Often, this is knowledge that some background condition connecting *e* to *p* doesn't obtain. In the drunk-friend example, knowing that the friend is drunk damages the credential of his testimony, because his testimony no longer makes the proposition in question more (or less) probable. Such knowledge allows (or even requires) you to ignore his testimony.

I have argued that knowledge that *e* is misleading is not in itself credential-damaging knowledge. However, it sometimes *induces* credential-damaging knowledge through inference-to-the-best-explanation. See this example:

---

[9] In reply, some defenders of principles similar to *K-norm* (e.g. Hawthorne and Stanley [2008: 588]) appeal to the phenomenon of 'pragmatic encroachment' within knowledge. I cannot enter into the issue here, but I do find the reply implausible, at least in this case. Even granting that practical stakes sometimes encroach within knowledge, it seems implausible that you no longer know that you have hands, simply by considering the bet.

[10] See Pettigrew [2013: 24—5] for a proof that matching credence with evidential probability (which is 'chance function given evidence' in his terms) is maximizing objective expected accuracy.

FALSE-REPORT

You missed a basketball game between team *A* and team *B* and you were wondering which team won. This morning you read *Guardian* (which you know is extremely reliable), and it says that (*p*) team *A* won the game. Call this evidence provided by *Guardian*'s report E. Based on E, you knew that *p*, and you became 0.9 confident in *p*. But now, you find that (*e*) *Times* reports that team *B* won the game. Gaining *e* doesn't deprive you of knowledge that *p*, because you know that, although *Times* is also reliable, it is much less reliable than *Guardian*. So, although you reduce your credence from 0.9 to 0.8 in the light of *e* (according to your background knowledge about the reliability of *Times*), you still know that *p*. So, you conclude that *Times*'s report is misleading with regard to *p*: it supports not-*p*, and yet not-*p* is false.

However, this conclusion makes you wonder how a reliable magazine like *Times* would make such a grave mistake. There must be an explanation. The best explanation, according to your background information, is that (C) *Times*'s report about the game is written by one of those few unreliable journalists from *Times* who make up stories through wishful thinking—in the past, when *Times* printed false reports, mostly it turned out that those reports were written by the few unreliable journalists. Then, you conclude that C must be true.

If inference-to-the-best-explanation is a way to expand knowledge (at least sometimes), you would know that C. This is credential-damaging knowledge. So, it seems that you can 'restore' your credence in *p* to 0.9, as if you have never read *Times*'s report at all.

Now, let me point out that this move cannot save *Entitlement*. *Entitlement* says that knowing that *e* is misleading *always* allows you to ignore it, whereas such knowledge only *sometimes* induces credential-damaging knowledge about *e*. For example, in SWIM, knowing that Philip wouldn't be drowned, and that therefore the evidence is misleading that he went swimming, doesn't induce knowledge that there is something wrong with the quality of the evidence.

A second point to note is that, even when knowledge that *e* is misleading does induce credential-damaging knowledge, you might not be allowed to ignore *e* *entirely*. The reason is that although I've been saying that credential-damaging knowledge allows you to ignore *e*, this is not exactly right, because the knowledge might not *entirely* destroy the credential of *e*.

In FALSE-REPORT, you come to know that C, but the probability of C is not 1. After all, you know that C by knowing that *p* and by performing an inference-to-the-best-explanation of the fact '*Times* reports not-*p*, and yet *p*.' Given that your knowledge that *p* doesn't come with probability 1, your knowledge that C also doesn't come with probability 1. So, there is a chance that *Times*'s report of not-*p* is not written by an unreliable journalist. If so, you cannot ignore the report in the sense of not revising credence even a little in the light of its report. That is, you cannot restore your credence to the level that you had before learning *Times*'s report, acting as if you have never learned of the report. Rather, you should take the report into account *to some degree*. And the degree should be a function of the probability of C. The more probable that C is, the less the degree to which you should take *Times*'s report into account. So, it seems that your credence in *p* after learning *Times*'s report is given by the following total probability principle:

$$\Pr(p/\text{B\&E\&}e) =$$
$$\Pr(p/\text{B\&E\&}e\text{\& C})\Pr(\text{C/B\&E\&}e) + \Pr(p/\text{B\&E\&}e\text{\&not-C})\Pr(\text{not-C/B\&E\&}e)$$

Here, 'B' refers to the propositions that figure in your inference-to-the-best-explanation. For example, it includes the proposition 'When *Times* printed false reports, mostly it

has turned out that those reports were written by the few unreliable journalists.' This equation shows that you can ignore e entirely as if you have never learned e, only if your credential-damaging knowledge that C comes with probability 1. Otherwise, your credence in p would be some value between Pr(p/B&E&e&C) and Pr(p/B&E&e&not-C), an average weighted by the probability of C.

So, the conclusions are as follows. (1) Although knowledge that e is misleading sometimes induces credential-damaging knowledge, it doesn't always; and (2) even when it does induce credential-damaging knowledge, it doesn't mean that you could entirely ignore e by not revising your belief or credence at all. The degree to which you could ignore e depends on how confident you should be that the credential of e is damaged.

7. Conclusion

Let's take stock. I have argued that the dogmatism puzzle can only be solved in a defeasibilist framework of knowledge. In this framework, I have explained why *Entitlement* is false. It appears true because knowledge that evidence e is misleading sounds like credential-damaging knowledge about e. I've argued that this is an illusion: such knowledge doesn't always undermine the credential of e, because e can still make p more (or less) likely to be true and therefore should still be respected. Although respecting such evidence is knowingly moving further away from the truth, it is striking the best balance between expected gain and loss. Finally, I have explained in what kind of cases knowing that e is misleading entitles one to ignore it, and I have argued that, even in those cases, one may ignore e not entirely but only to some degree. Hopefully, my discussion of the dogmatism puzzle makes it clearer in what sense rational believers rely on evidence and when it is that we can ignore evidence.[11]

References

Baumann, Peter 2013. Knowledge and Dogmatism, *The Philosophical Quarterly* 63/250: 1−19.
Cargile, James 1995. Justification and Misleading Defeaters, *Analysis* 55/3: 216−20.
Harman, Gilbert 1973. *Thought*, Princeton: Princeton University Press.
Hawthorne, John 2004. *Knowledge and Lotteries*. Oxford: Clarendon Press.
Hawthorne, John and Jason Stanley 2008. Knowledge and Action, *The Journal of Philosophy* 105/10: 571−90.
Lasonen-Aarnio, Maria 2014. The Dogmatism Puzzle, *Australasian Journal of Philosophy* 92/3: 417−32.
Neta, Ram 2009. Defeating the Dogma of Defeasibility, in *Williamson on Knowledge*, ed. Patrick Greenough and Duncan Pritchard, Oxford: Oxford University Press: 161−82.
Pettigrew, Richard 2013. A New Epistemic Utility Argument for the Principal Principle, *Episteme* 10/1: 19−35.
Sharon, Assaf and Levi Spectre 2010. Dogmatism Repuzzled, *Philosophical Studies* 148/2: 307−21.
Sorensen, Roy 1988. Dogmatism, Junk Knowledge, and Conditionals, *The Philosophical Quarterly* 38/153: 433−54.
Williamson, Timothy 2000. *Knowledge and Its Limits*, Oxford: Clarendon Press.