

## On Being Debased

**ABSTRACT:** A standard form of skeptical scenario, in the tradition of Descartes' evil demon, raises the prospect that our sensory experiences are deceptive. A less familiar and importantly different kind of skeptical scenario raises the prospect that our beliefs have been *debased* (Schaffer 2010). This paper provides a new and improved way of resisting this latter kind of debasing skepticism. Along the way, I explore how the debasing demon scenario connects with some potentially controversial epistemological principles and clear up various neglected or misunderstood points concerning debasing skepticism.

**KEYWORDS:** Epistemology, Skepticism, The Basing Relation, Rationality

### 1. Introduction

Skeptical scenarios are stories that a skeptic spins in order to support the surprising, revisionary conclusion that we lack knowledge that we normally, intuitively take ourselves to have. Perhaps the most famous such skeptical scenario is Descartes' story of an evil demon<sup>1</sup>, who can manipulate our perceptual experiences and create deceptive appearances of our external surroundings. Almost as famous – and much more ancient – is the story 'Dreams of Being a Butterfly' from the second inner chapter of *The Zhuangzhi*, which suggests we cannot know whether we are dreaming or awake. More recent variants of such stories might appeal to brains-in-vats or virtual reality machines<sup>2</sup>. These sorts of skeptical scenarios will presumably be familiar to anyone who has taken an Intro to Philosophy course – or to anyone who has watched 'The Matrix'.

An importantly different and less familiar kind of skeptical scenario is based on the idea of '*debasing*' – see Schaffer (2010). Whereas, the traditional Cartesian Deceiving Demon prevents one from gaining knowledge by providing false appearances or misleading evidence that lead one into forming false beliefs, a *Debasing Demon*, in contrast, prevents knowledge by causing one to form beliefs that are not properly based on one's evidence (propositional justifiers). Although it still seems that one's beliefs are properly based, in fact they are formed via some improper, irrational basis, e.g. superstition, or wishful thinking, or perhaps they are just based on something totally irrelevant. So, the Debasing Demon does not promote false beliefs<sup>3</sup>, rather it promotes *debased* beliefs. The feature of the Debasing Demon scenario that Schaffer himself most emphasized as contrasting with the Cartesian Deceiving Demon is that the Debasing Demon (allegedly) threatens universal skepticism. The misleading appearances conjured by a Cartesian Deceiving Demon may threaten all of our knowledge about the external world, but they seem to leave intact knowledge about our own current conscious

---

<sup>1</sup> In the 1<sup>st</sup> meditation Descartes considers both the idea of a deceiving god and also the idea of a *genium malignum* (or *mauvais génie*), which has sometimes been translated as 'evil genius' and sometimes as 'evil demon'. In the 2<sup>nd</sup> meditation, he writes of 'some deceiver or other, very powerful and very cunning, who ever employs his ingenuity in deceiving me' (Descartes 1984[1641], 2:16)

<sup>2</sup> See, for example, the discussions in Putnam 1981 or Chalmers 2022.

<sup>3</sup> Though notice that the Debasing Demon would tend to produce false beliefs specifically about the basing of one's own beliefs. See the discussion in section 3, below.

experiences and perhaps also simple a priori reasoning<sup>4</sup>. However, if it is a requirement on all knowledge that it is properly based on a justifying reason, in a debasing demon scenario one would lack any knowledge – even knowledge that one currently exists, or knowledge that either one exists or one does not exist, etc.

‘The debasing demon thus stands as a universal demon. In this respect, the debasing demon seems more nightmarish than the deceiving demon. There may be limits as to how far we can be deceived (for instance if Descartes is right then we cannot be deceived about our own present existence), but there are no limits whatsoever to how far we can be debased.’ (Schaffer, 2010, 234)

It is unsurprising that this genuinely novel kind of skepticism generated a flurry of responses and discussions by various other epistemologists – e.g. Brueckner (2011), Ballanytne & Evans (2013) Conee (2015), Bondy & Carter (2018), Smithies (2019), McCain (2020), Cunningham (2021). However, despite all this attention, I think that there has not yet been a sufficiently convincing response to the skeptical threat raised by the debasing demon. The main aim for this paper then will be to provide a new and improved way of resisting the debasing skeptic. It is a mark of how rich the Debasing Demon scenario is that it raises various interesting and difficult questions about higher-order evidence, defeat, and justification. And so a secondary aim will be to show how the Debasing Demon scenario connects with three potentially controversial epistemic principles – though without myself endorsing or denying any of these principles.

The plan for the paper is as follows: in section 2 I discuss the nature of the sceptical threat posed by the mere possibility of the debasing demon scenario. In section 3 I clear up some misunderstandings and discuss some complexities concerning one’s epistemic situation if the debasing scenario actually obtains or if one has positive evidence that it actually obtains. In section 4 I discuss and criticize Bondy & Carter’s (2018) reply to the debasing skeptic. In section 5 I argue for a new response to the debasing skeptic. Section 6 provides a short summary and conclusion.

## **2. How can the mere possibility of a debasing demon threaten Skepticism?**

Anthony Brueckner (2011) makes the objection that for the *mere possibility* of the Debasing Demon to pose a universal skeptical threat, a skeptical argument would need to rely on a version of the KK principle. And in fact Schaffer (2010) acknowledges as much in a footnote (see below). In this section I will consider whether and how the mere possibility of the debasing demon scenario might still pose a significant skeptical threat even without a KK principle. Of course, if I *actually* am the victim of a debasing demon then, by hypothesis, I have no knowledge whatever. But there is not the slightest shred of positive reason to think that I am actually the victim of any such debasement. Generally, the point of raising skeptical scenarios is not to positively convince us that the skeptical scenario really does obtain. Rather, the skeptic’s aim is to get us to admit that we have no rational/justified basis to rule out the skeptical hypothesis – or perhaps the slightly stronger claim that we have no rational/justified basis to *prefer* the non-skeptical hypothesis to the skeptical hypothesis. And

---

<sup>4</sup> As Schaffer (2010) notes, Descartes seems to vacillate a bit concerning whether simple *a priori* knowledge would be threatened by his method of doubt.

so (it is supposed to follow) we cannot after all know many of the things we take ourselves to know. In other words, it is the mere possibility of the skeptical scenario that is meant to pose a threat to knowledge, whether or not the scenario is actual. So how is this supposed to work in the case of the debasing demon? Here is the argument as formulated by both Schaffer and Brueckner:

- (1) If I know that p then I have a properly based belief that p. [Basing is required for K]
- (2) If I know that p then I can know that I know that p. [weak form of KK principle]
- (3) If I can know that I know that p then I can know that I have a properly based belief that p. [from 1 by Closure]
- (4) If I know that p then I can know that I have a properly based belief that p. [from 2, 3]
- (5) I cannot know that I have a properly based belief that p. [Debasing hypothesis is epistemically possible.]
- (6) I do not know that p. [from 4, 5 by contraposition]

Premise 1 is meant to be an uncontroversial requirement on knowledge<sup>5</sup>. Premise 3 comes from Premise 1 plus the Closure of what can be known over knowable entailment. Premise 4 follows from 2 and 3. Premise 5 is supposed to flow from the fact that you cannot rule out the possibility of the Debasing Demon. But premise 2 is a controversial and non-obvious principle that has generated much debate and disagreement<sup>6</sup>. Brueckner concludes that Schaffer’s debasing skepticism ‘has the drawback of requiring a strong and implausible epistemic principle’ (Brueckner, 2011, 297). And Schaffer himself says:

“there is also a respect in which the debasing demon might seem less threatening, for perhaps the most natural route from the possibility of debasement to actual ignorance goes through a questionable premiss from knowing to being in an evidential position to know that one knows...” (Schaffer, 2010, footnote 7, p234)

In order to pose a skeptical threat must the debasing skeptic rely on a KK principle? The argument above could be simplified by relying directly on an alternative general epistemic principle that: if I know that p then I can know that I have a properly based belief that p. Then one could simply reason as follows<sup>7</sup>:

- (1\*) If I know that p then I can know that I have a properly based belief that p.
- (2\*) I cannot know that I have a properly based belief that p.
- (3\*) I do not know that p.

---

<sup>5</sup> Though see Wedgwood (2022) for a denial that we should think of doxastic justification in terms of a basing relation.

<sup>6</sup> This is, of course, a weaker principle than the most basic  $Kp \rightarrow K(Kp)$  principle, claiming only that  $Kp$  entails being *in a position to*  $K(Kp)$ . Williamson (2000) provides an influential argument against the weaker principle. See also Dorst (2019). For some recent defences of this KK principle see Greco (2014), Das & Salow (2018), Goodman & Salow 2018.

<sup>7</sup> This formulation of a debasing-skeptical argument is similar to Conee’s (2015) formulation. In particular, premise (1\*) is very close to being a contraposed version of what Conee labels the ‘Knowability Premise: If none of us is able to know that our beliefs are properly based, then we lack knowledge’, though Conee’s version is put in terms of what we can know rather than what I can know. Conee’s response to the debasing skeptic is effectively to deny (2\*) and argue that we can fallibly know our beliefs are properly based. See section 5, below, for a brief discussion of Conee’s position.

Premise (1\*) might seem somewhat weaker and so easier to defend than the KK principle (premise (2) of the first argument) insofar as it only requires that we can know about something internal to our own minds – that the proper basing relation (whatever that is exactly) obtains. Nevertheless, premise (1\*) is still far from obvious. For a start there will be issues with cognitively unsophisticated subjects, such as children and animals, to whom we are intuitively happy to ascribe knowledge yet who lack any concept of the basing relation<sup>8</sup>. But even with more sophisticated subjects it is just not obvious that a lack of knowledge that a belief is properly based would always amount to a defeater for that belief counting as knowledge. For example: I know that  $2+2=4$ . Suppose I now get some misleading evidence that my beliefs may be debased. Let's grant that I am no longer in a position to know that my belief that  $2+2=4$  is properly based. It is not obvious that this would entail that I no longer know that  $2+2=4$ . You might think instead that so long as my belief that  $2+2=4$  is in fact properly based, then I can still count as knowing that  $2+2=4$  despite now being unsure about the basing relation<sup>9</sup>. I am not suggesting that this is an immediately decisive or persuasive objection to (1\*), no doubt more could be said on both sides. The point is simply that (1\*) – that knowing entails being in a position to know about the basing relation – is still a highly questionable premise to rely on when arguing for the radical skeptical conclusion that we lack knowledge entirely. Cunningham (2021) offers an alternative way of moving from the mere possibility of the debasing scenario to a universal skeptical conclusion that does not rely any KK principle, but instead appeals to a 'justificationist' construal of the basing relation together with various theses about performing basic actions and our knowledge of and evidence for such actions. Space does not permit any detailed discussion or evaluation of Cunningham's interesting but complex version of debasing skepticism here. But of course, as Cunningham himself explicitly notes, rather than accepting the extremely counter-intuitive skeptical conclusion that we lack knowledge entirely, we might instead treat his argument as a *reductio* of the various philosophical theses that it relies on.

The first thing I would like to point out is that even if we set aside the KK principle, or premise (1\*), or any of the potentially questionable assumptions that Cunningham relies on, the mere possibility of the debasing demon can still straightforwardly threaten skepticism about any *higher-order knowledge* (knowledge about one's own knowledge, knowledge about one's own doxastic justification). This falls short of the universal skepticism that Schaffer advertised, but it would still be a pretty interesting and potentially radical kind of skepticism; something that, if true, would be deeply surprising and counterintuitive<sup>10</sup>. To see how, consider a standard way of laying out the threat posed by the mere possibility of the Cartesian Deceiving Demon:

---

<sup>8</sup> This kind of worry could perhaps be avoided if we can make sense of a subject being 'in a position to know' a proposition that they lack the conceptual sophistication to entertain or understand.

<sup>9</sup> After all, I can still perfectly well access and appreciate my reasons – the propositional justifiers – for my belief that  $2+2=4$ .

<sup>10</sup> Silva (2019), citing a lecture by John Hawthorne, points out that everyday knowledge of our own knowledge should be as abundant and easy as knowing what other people know. E.g. if I see you pick up the keys from the table, I will usually come to know that you know where the keys were located. Likewise then, if I were to pick up the keys from the table, I should normally be able to know that I know where the keys were located. So denying this sort of everyday higher-order knowledge would seem to threaten our everyday knowledge of other people's mental states.

(1#) We have no rational basis to prefer a normal, non-skeptical hypothesis (here are two hands, there is a laptop on the table, etc.) to the Cartesian Demon hypothesis (there is no hand or laptop in front of me, just a demon-caused appearance).

(2#) If we have no rational basis to favour a hypothesis, H, over the hypothesis that not-H then we don't know H. (Nor do we know not-H.)

(3#) SO: We don't know the non-skeptical hypothesis (here are two hands, there is a laptop in front of me, etc.)

To be clear: I am not endorsing this skeptical argument. It is just meant to be an object of comparison. There are, of course, ways one might try to resist. Most obviously one might reject the initial premise (1#) by insisting that although the skeptical hypothesis is consistent with our sensory evidence, nevertheless it can be rational to prefer the non-skeptical hypothesis – for example by appeal to abduction and/or to considerations of simplicity<sup>11</sup>. But I take it that something like this line of thought is one standard formulation of the idea that the mere possibility of the Cartesian Deceiving Demon scenario poses a skeptical threat. And the only general epistemic principle that the argument relies on is premise (2#), which is surely much less controversial than (2)'s KK principle, or (1\*), or a justificationist view of the basing relation. Now compare:

(1##) We have no rational basis to prefer the normal non-debased hypothesis (in which many of our beliefs amount to knowledge) to the Debasing Demon hypothesis (in which we do not know anything).

(2##) If we have no rational basis to favour a hypothesis, H, over the hypothesis that not-H then we don't know H. (Nor do we know not-H.)

(3##) SO: We don't know that any of our beliefs amount to knowledge.

The form of argument is the same - the general epistemic principle in (2##) is identical to (2#). But here the conclusion is that we have no knowledge of our own knowledge<sup>12</sup>. This is still a pretty radical conclusion to reach from the mere possibility of the Debasing Demon

---

<sup>11</sup> This is effectively the line taken by Russell (1912). See Vogel (1990), Feldman & Conee (2004) for useful discussion. Conee (2015) discusses this sort of anti-skeptical strategy with specific reference to the debasing scenario.

<sup>12</sup> We could alternatively think in terms of closure-based formulations of skepticism. The standard closure-based argument against external world knowledge runs as follows ('Two Hands' abbreviates the proposition that there are two hands in front of me, 'BIV' abbreviates the proposition that I am a brain in a vat):

- |   |                         |
|---|-------------------------|
| (1) $[K(\text{Two Hands}) \ \& \ K(\text{Two Hands} \rightarrow \neg\text{BIV})] \rightarrow K(\neg\text{BIV})$ | Closure                 |
| (2) $\neg K(\neg\text{BIV})$  | BIV cannot be ruled out |
| (3) $\neg[K(\text{Two Hands}) \ \& \ K(\text{Two Hands} \rightarrow \neg\text{BIV})]$                           | From 1, 2               |
| (4) $K(\text{Two Hands} \rightarrow \neg\text{BIV})$  | A priori Knowledge      |
| (5) $\neg K(\text{Two Hands})$  | From 3, 4               |

We can run a parallel closure-based argument against meta-knowledge appealing to the possibility of debasement:

- |   |                             |
|---|-----------------------------|
| (1) $[K(Kp) \ \& \ K(Kp \rightarrow \neg\text{Debased})] \rightarrow K(\neg\text{Debased})$ | Closure                     |
| (2) $\neg K(\neg\text{Debased})$  | Debased cannot be ruled out |
| (3) $\neg[K(Kp) \ \& \ K(Kp \rightarrow \neg\text{Debased})]$                               | From 1, 2                   |
| (4) $K(Kp \rightarrow \neg\text{Debased})$  | A priori Knowledge          |
| (5) $\neg K(Kp)$  | From 3, 4                   |

scenario. Whilst it is not full universal skepticism, ruling out knowledge that  $p$  for any  $p$ , the conclusion would rule out knowing you know that  $p$  for any  $p$ . So even without relying on any extra, potentially controversial epistemic principles, the mere possibility of the Debasing Demon still seems to provide a skeptical threat to any 2<sup>nd</sup>-order knowledge, that is of the same form – and so *prima facie* of the same strength – as the threat provided by the mere possibility of the Cartesian demon against 1<sup>st</sup>-order external world knowledge<sup>13</sup>.

A second thing that is important to notice is that a debasing demon could equally well debase any of the other doxastic attitudes, not just full belief. A debasing demon could make it the case that your credences (degrees of belief) are not properly based on your evidence, though they still seem to be. And likewise that your attitude of suspending judgement is not properly based on your evidence, though it still seems to be. This is important since a skeptic who raises the mere possibility of some skeptical scenario typically does so because they want to persuade us to give up some of our beliefs and to *suspend judgement instead*. E.g. a skeptic raises the possibility of the Cartesian demon in order to persuade us to give up our beliefs about the external world and to suspend judgement instead. But when it comes to a debasing demon scenario it seems that suspending would not be any better than belief, rationally speaking, given that a debasing demon can debase suspension too. Why should invoking the mere possibility of a debasing demon persuade us to suspend judgement if it is just as much of a possibility that this demon is debasing the attitude of suspending as debasing the attitude of belief? Now, to be clear: *if* a persuasive skeptical argument could be made for the claim that suspending judgement whether  $p$  is the propositionally justified attitude that we ought to take to  $p$ , this would be philosophically interesting *even if* the premises also implied that suspending would not be doxastically justified (i.e. would not be properly based). However, when it comes to normal, everyday propositions about the external world, the mere possibility of the debasing demon does *not* seem to provide an argument that suspending judgement is the propositionally justified attitude to take to such propositions. Recall the skeptical argument as formulated by Schaffer and Bruckner: this argument went from the claim that we are not in a position to know that our belief that  $p$  is properly based in order to conclude, via a KK principle, that we do not know that  $p$ . But this does not give us any reason to think that suspension is therefore the *propositionally justified* attitude to take towards  $p$ . On the contrary, all our evidence might still propositionally justify that  $p$  is true<sup>14</sup>.

The foregoing discussion points us towards the following potentially controversial principle:

- **Potentially Controversial 1 (PC1):** If you have no rational basis to prefer the hypothesis that forming some mental state,  $M$ , would be rational to the hypothesis that forming  $M$  would be not be rational, then you rationally ought not go ahead and form  $M$ .

---

<sup>13</sup> Though again, see Conee (2015) who argues that we actually have *better* fallibilist grounds to reject what I am labeling (1##) in the debasing demon scenario than to reject (1#) in the standard Cartesian demon scenario.

<sup>14</sup> It is only when the proposition in question,  $p$ , specifically concerns having properly based beliefs that the mere possibility of undetectable debasing could provide an argument that the propositionally justified attitude towards  $p$  is suspension of judgement.

Whether to accept this principle will depend on difficult and controversial issues concerning how a subject's evidence (or lack thereof) about what would be rational bears on what is actually rational for the subject. I will not attempt to say anything for or against PC1 here. But if PC1 is accepted we get a simple argument that the mere possibility of the debasing demon entails that it is not rationally permitted to believe anything:

- (I) We have no rational basis to prefer the normal non-debased hypothesis to the Debasing Demon hypothesis in which any beliefs we form are debased, hence irrational.
- (II) If you have no rational basis to prefer the hypothesis that forming some mental state, M, would be rational to the hypothesis that forming M would be not be rational, then you rationally ought not go ahead and form M. (PC1)
- (III) SO: it is not rationally permitted to form any beliefs whatsoever.

If this skeptical argument is correct it would mean that the mere possibility of the debasing demon scenario would not only make knowledge impossible, it would also make rational belief impossible. Now let's take into account that a debasing demon could debase any doxastic attitude<sup>15</sup>, including suspension, not just belief – let's call this the Debasing Demon+ hypothesis. This means that we can run an argument, again relying on PC1, from the mere possibility of the debasing demon to the conclusion that it is not rationally permitted to form any doxastic attitude to anything – i.e. that we are caught in a universal rational dilemma.

- (I\*) We have no rational basis to prefer the normal non-debased hypothesis to the Debasing Demon+ hypothesis in which any doxastic attitudes we form are debased, hence irrational.
- (II\*) If you have no rational basis to prefer the hypothesis that forming some mental state, M, would be rational to the hypothesis that forming M would be not be rational, then you rationally ought not go ahead and form M. (PC1)
- (III\*) SO: it is not rationally permitted to form any doxastic attitude to any proposition whatsoever.

Given that a debasing skeptic will presumably endorse the first premises of these arguments, it seems that any debasing skeptic who wishes to to avoid the conclusions must deny PC1<sup>16</sup>. But notice that something like PC1 would seem to be needed by a debasing skeptic if the skeptic wants to move from the premise that a debasing scenario is consistent with all our evidence (i.e. for all we can tell we are currently debased) to the conclusion that continuing to form beliefs as usual would not be rationally permitted. On the other hand, if PC1 is accepted then, by the debasing skeptic's own lights, we are all caught in a universal rational dilemma where no doxastic attitude to any proposition would ever be rational. Since for any doxastic attitude to any proposition there could be a possible scenario of undetectable debasing that is consistent with all our evidence. This would still be a skeptical result, but it would be one that we could do absolutely nothing rational about in response, since any doxastic attitude we form would be irrational. And so it seems there would be no rational pressure to do anything

---

<sup>15</sup> Indeed a debasing demon could debase any attitude that rationally requires proper basing on reasons, which would also include non-doxastic attitudes such as intentions.

<sup>16</sup> Some theorists, such as Titelbaum (2015), Smithies (2019), Tal (2020), maintain that we always have indefeasible *a priori* justification to believe truths about which attitudes would be rational/irrational. On this kind of view the antecedent in PC1 would presumably always be false. Which would make PC1 trivially true but uninteresting.

*differently* so far as forming beliefs and other doxastic attitudes is concerned. Notice also that the skeptic would then face the issue that if no doxastic attitudes are ever rational then nobody, including the skeptic herself, can ever rationally believe the skeptic's own premises or conclusions. Nor could we ever rationally believe the consequence that no doxastic attitudes are ever rational. Nor could we rationally disbelieve or suspend either. However, one could still have a true but irrational belief that no doxastic attitudes are ever rational.

In summary: (i) even without relying on any potentially controversial general epistemic principles (such as KK) that go beyond what are required for normal Cartesian skeptical arguments, the mere possibility of the debasing demon still poses a sceptical threat to any higher-order knowledge. (ii) The mere possibility of the debasing demon scenario cannot be raised as a reason to suspend judgement rather than believe, since suspension is just as vulnerable to demonic debasing as belief. (iii) If a potentially controversial principle, PC1, is accepted, then the mere possibility of a debasing demon would not only threaten all knowledge, it would threaten to make it irrational to form any doxastic attitude whatever.

### 3. What if the Debasing Demon Scenario actually obtained?

Having considered the skeptical threat posed by the mere possibility of the debasing demon scenario, let's turn now to considering the subject's epistemic position if the debasing demon scenario *actually obtained* – and also if the subject had *positive evidence* that the scenario actually obtains.

Here is a quote from Ballantyne & Evans (2013), in which they consider the debasing demon scenario as actual:

'...pick at random a first-order belief of yours that  $p$ . As we see it, the possibility that you are debased is *irrelevant* to your belief that  $p$ . For the possibility that you are debased is a possibility in which  $p$  is still true. Even if you were to discover that you in fact were debased, that should not change your answer to the question whether  $p$ : for you would know (by the definition of a debasing scenario) that  $p$  was true – you would know that you were in a  $p$ -world.' (Ballantyne & Evans, 2013, 555)

Let's distinguish between the following two cases<sup>17</sup>:

- (i) You have been a victim of the debasing demon up until now, but now you are no longer being debased.
- (ii) You have been and still currently are being debased.

If you are still currently being debased, as in case (ii), then by hypothesis you cannot know anything at all. And so in particular you could not 'know by the definition of a debasing scenario that  $p$  was true' nor could you 'know that you were in a  $p$  world'. (Nor would you know that you have propositional justification/evidence for your beliefs, etc.) Indeed, if the notion of 'discovering' something is understood as coming to know something, then a current victim of the debasing demon simply cannot discover that they are being debased, at best they could have an irrational, unjustified but true belief that they are currently debased. (Of course

---

<sup>17</sup> Ballantyne and Evans switch between the phrases 'you are debased' and 'you in fact were debased', which seem to correspond to these two different cases.



one can discover/know that one was previously, but is no longer, debased.) So if you are actually and currently being debased and someone asks you the question: ‘whether p?’ you certainly cannot know that the answer is p, even though you could still utter the answer ‘p’ and you could still have a true but debased belief that p.

Somewhat similar issues arise when we consider how a subject should respond to evidence which positively indicates that they are the victim of a debasing demon. Here is a passage from Declan Smithies’ excellent recent book:

‘Suppose you know that your beliefs are properly based on evidence, but then you acquire misleading evidence that you’re the victim of the debasing demon. **What should you do? Nothing at all. There is no rational pressure to revise your beliefs in any way.** The rational response is simply to maintain the justified beliefs you already have on the basis on which you already hold them. After all, **you know that your beliefs are supported by the evidence, although you don’t know that they are properly based on the evidence, since you have misleading evidence to the contrary...**’ (Smithies, 2019, 336, bold type added)

A minor criticism of this passage concerns meta-beliefs – e.g. a belief that some belief of yours is properly based, or is doxastically justified, or is knowledge. In the case of these beliefs the misleading evidence of debasement does create ‘rational pressure’ to revise your belief, which is no longer even propositionally justified. So there would be ‘rational pressure’ to revise any of your meta-beliefs about your own knowledge and your own doxastic justification. Smithies’ claim that you should do ‘nothing at all’ thus cannot be correct when it comes to such meta-beliefs. But if we set aside such higher-order beliefs about knowledge or the basing relation etc. and restrict ourselves to 1<sup>st</sup>-order beliefs about the external world, then I do agree with Smithies that in response to evidence that you have been debased ‘there is no rational pressure to revise your beliefs’, at least in the sense of *changing* what you believe<sup>18</sup>. However, it seems that there would still be some kind of rational pressure to try to, as it were, *re-base* – i.e. reform or re-affirm – one’s beliefs so that they are properly based on one’s evidence. After all, suppose some super-reliable, super-intelligent expert source tells you: ‘Though it may appear otherwise, your belief that p was irrationally formed. But in fact you already possess all the evidence you need to form a justified/rational belief that p.’ It seems that even though you have been assured that your belief has propositional justification, there is still some rational pressure to review your evidence and re-form or re-assert your belief that p so that it is properly formed, hence doxastically justified. And I think this raises an interesting potential puzzle. Once more let’s explicitly draw a distinction between the following two cases:

- (a) You get compelling evidence that: up till now you have been the victim of the debasing demon, but you are no longer being messed around with.
- (b) You get compelling evidence that: you have been *and still currently are* the victim of the debasing demon.

---

<sup>18</sup> Also setting aside, of course, non-skeptical beliefs such as ‘there are no malicious demons’, or, ‘I am not the victim of someone messing with my brain’, etc.

In case (b), where the evidence is that you are still being debased, it is at least doubtful whether it would be rationally permitted for you to respond to this evidence by forming the belief that you are currently being debased, since this very same evidence would also indicate that any *new* beliefs you form will end up being de-based, hence irrational and not knowledge. So it is at least unclear that you would be rationally permitted to form or re-form any new beliefs whatsoever. This relies on a second potentially controversial principle:

- **Potentially Controversial 2 (PC2):** If you have sufficiently strong evidence that believing that *p* would not be rational, then it is not rationally permitted for you to believe that *p*.

Again, whether to accept this principle will be tied up with difficult and controversial issues about higher-order evidence and misleading evidence about rational requirements, which I cannot attempt to settle here<sup>19</sup>. There is also a distinction between assessing the rationality of a belief-forming process versus assessing the rationality of the resultant belief<sup>20</sup>, which could well be important for assessing PC2. But setting these issues aside, if we were to grant PC2, then, once more, the threat of a dilemma arises. For on the one hand, really strong evidence that you are being debased would not rationally permit believing that you are currently being debased. But on the other hand it also does not seem rationally permissible to just entirely ignore extremely strong evidence that one has been debased<sup>21</sup>. So it seems that it would not be rationally permitted to suspend judgment as to whether one has been debased either. Evidence that you are currently being debased would thus seem to threaten a kind of epistemic dilemma concerning the proposition that you are currently being debased. More generally, any situation in which your total evidence propositionally justifies believing that *p* but also indicates that you are currently debased is one in which your total evidence indicates that *p* is true but also that a belief that *p* would not be rational if formed. So this total evidence is pointing towards the truth of the following Moorean-flavoured conjunction: *p* & I ought not believe that *p*<sup>22</sup>. This total evidence also still supports the conjunction: *p* & it is rationally permitted to believe that *p* *on a proper basis*. But the evidence indicates that if I form a belief that *p* it will not be properly based, hence not rational. Assuming PC2 is correct, if I were to go ahead and believe *p* this belief would not be rational, *even if I am not in fact being debased*. For according to PC2, so long as my evidence indicates that I am being debased, so indicates that any belief that *p* I form will be irrational, then it really is not rationally permitted to believe that *p* (no matter how strong one's evidence for *p*).

In summary: (i) if you actually are a current victim of the debasing demon then you cannot know that you are, no matter how strong the evidence of debasing. (ii) Even if you get very strong evidence that you are currently being debased, it is at least doubtful whether you can rationally believe that you are currently being debased – since this evidence indicates that any

---

<sup>19</sup> The recent literature on higher-order evidence and defeat is enormous and seems as yet to have reached little consensus. See for example: Titelbaum (2015), Christensen (2010), Horowitz (2014), Lasonen-Aarnio (2014, 2020), Pryor (2018), Weatherson (2019).

<sup>20</sup> See Silva (2022) for helpful discussion of this distinction.

<sup>21</sup> This could perhaps be resisted if one endorses the idea that there are 'no positive epistemic duties' and so it is always permissible to suspend judgement whether *p*, no matter how strong one's evidence for *p*. See e.g. Nelson (2010).

<sup>22</sup> See Lasonen-Aarnio (2014, 2020) for defence of the idea that it can sometimes be rational to believe these kinds of conjunctions.

beliefs you form will be debased hence irrational. But given the strength of the evidence, it is also doubtful that it would be rational to suspend judgement either. So very strong evidence that you are currently being debased threatens to create a dilemma where no attitude would be rational. The issues here turned on another potentially controversial principle, PC2.

#### 4. Would the ‘Debasing Demon’ even be debasing?

Let’s turn now to considering how the threat of the Debasing Demon Scenario might be resisted. In Schaffer’s original presentation of the debasing demon, he is clear that the demon is supposed to *take control* of the subject’s belief-forming mechanisms:

“The debasing demon ... **throws her victims into the belief state** on an improper basis, while leaving them with the impression as if they had proceeded properly. So for instance, the debasing demon might **force me into believing** that I have hands on the basis of a blind guess or mere wishful thinking, while leaving me with the impression as if I had come to this belief on the basis of the visual evidence.” (Schaffer, 2010, p231, bold-type added)

One anti-skeptical strategy here is to suggest that this interfering with and taking control of the subject’s cognitive processes by the demon would undermine the skeptical power of the debasing scenario. A version of this strategy is pursued in a recent paper by Bondy & Carter (2018). Whilst I think that there is something importantly right about this general idea, in this section I will argue that Bondy & Carter do not provide a persuasive development of it. Then in the next section I will present what I take to be a better way of resisting the threat of the debasing demon.

Bondy & Carter understand the debasing demon scenario as requiring the following three conditions to hold<sup>23</sup>:

- “(i) there is a possible world where a demon causes some subject *S* to believe *p* on the basis of a bad reason *R1* at time *t*.
- (ii) at a later time *t2* the demon makes *S* think that his reason for believing *p* was a different good reason *R2*.
- (iii) at *t2*, it remains the case that *S* holds his belief only on the basis of *R1*, and not on the basis of *R2*.” (Bondy & Carter, 2018, 210)

Bondy & Carter then argue that on either of the two main views on the basing relation these three conditions cannot all be fulfilled. And so, they maintain, the subject in the supposed debasing scenario would not actually end up with a debased belief. What are these two main

---

<sup>23</sup> It is unclear why Bondy & Carter think that a debasing scenario *requires* that the demon causes the subject to be mistaken about the basis for her belief at a later time, *t2*, than when the demon causes the subject to form the debased belief (at *t1*). Though of course it is *possible* for a debasing scenario to play out this way, the idea that creating the misleading appearance of proper basing occurs at some later time than the debased belief-formation does not seem to be part of Schaffer’s original presentation of the case (see the quote above), nor does it seem to be required. Why couldn’t the demon cause me to base my belief on a bad reason *R1* whilst *simultaneously* producing the false appearance that I am basing it on a good reason *R2*?

views? Well, leaving out many details and possible complications, the two main kinds of views about basing are:

- CAUSAL THEORIES OF BASING: S's belief that p is based on reason R iff S's belief that p is *caused in the right way* by R.
- DOXASTIC THEORIES OF BASING: S's belief that p is based on Reason R if/only if/iff S believes (meta-belief) that R is S's reason to believe that p. (Or that S is a good reason to believe that p.)

There are various arguments for/against both of these broad views of the basing relation<sup>24</sup>, which we can ignore for present purposes as Bondy & Carter's point is that on *neither* account would the subject end up debased. Let's consider the causal theory first.

Very briefly, it is generally accepted that any plausible version of the causal theory of basing requires that the reason R causes the belief that p *in the right way*. This is to rule out *deviant* causal chains as when, e.g. my belief that there is a burglar in the house causes me to drop my cup of tea, which smashes on the ground and scalds my leg, which causes me to believe my leg hurts. The former belief is a crucial cause of the latter belief, but it is not *my reason* for forming the latter belief – this is not an instance of the basing relation. It is actually very tricky to specify in a non-circular way what makes a causal chain deviant or non-deviant, but it would be a needless distraction here to get into the details of examples and counter-examples. The point is just that according to Bondy & Carter if a demon intervenes in the subject's mental processes so as to make some bad reason R1 cause the belief that p then this would plausibly (so they claim) be a case of a *deviant* causal chain and so the subject's belief would not after all be based on bad reason R1. Let's just grant, for the sake of argument, that Bondy & Carter are correct that in the debasing demon scenario the causal chain involving the bad reason R1 would be deviant and so the subject's belief would not be based on reason R1. This does not yet defuse the threat of the debasing demon scenario. For then, by the lights of the causal theory of basing, it would simply mean that the subject's belief (caused by the demon in a way that is by assumption deviant) is *not based on any reason at all*. And this would presumably be just as ruinous to knowledge as being based on a bad reason<sup>25</sup>. So long as we allow that the debasing demon causes the subject to possess a belief that is not properly based on a good reason, then the subject will lack knowledge. It really doesn't make any difference to this outcome if the demon causes the subject to possess a belief that is based on a bad reason or if the demon causes the subject to possess a belief that is not based on anything at all.

Let's now consider what Bondy & Carter have to say about the other, doxastic account of the basing relation. Bondy & Carter's core point here is that if we understand the doxastic view as stating either a necessary or a sufficient condition for a belief to be based on a reason R, it turns out that debasing demon scenario is not really possible.

---

<sup>24</sup> Two helpful overviews of work on the basing relation are Korcz (1997) and Sylvan (2016). Carter & Bondy (eds.) (2020) is a recent collection of essays devoted to the basing relation. See also: Neta (2019), Silva Jr. (2015), McCain (2012).

<sup>25</sup> I am here in complete agreement with Cunningham (2021).

Understood as a *necessary* condition the doxastic view claims that: if my belief that *p* is based on reason *R*, then I must believe that *R* is my reason to believe that *p*. And here Bondy & Carter claim that in the alleged debasing scenario the subject would lack any such meta-belief that the bad reason *R1* is her reason to believe that *p*. After all, the subject is meant to be totally unaware of the involvement of *R1*, which is supposed to have been somehow hidden from conscious awareness by the demon. However, as it stands Bondy & Carter's objection is too quick. For it is presumably possible for the demon, as part of the debasing process, to instill the required meta-belief in the subject (that the bad reason *R1* is her reason/a good reason to believe that *p*) but to keep it hidden from conscious introspective awareness. If we are considering the possibility that a demon can create unconscious bad-basing, why not allow that the demon can create unconscious meta-beliefs about the basing? Now perhaps more could and should be said here. For example, the idea that in the debasing scenario we would have these sorts of extra unconscious beliefs might, depending on one's preferred conception of evidence, tend to undermine the idea that this skeptical scenario is consistent with our total evidence, since as a matter of fact we do not have these extra unconscious beliefs. Or, perhaps an amended version of the doxastic view of basing could include the requirement that the relevant meta-belief about basing must be introspectively accessible. But pursuing these sorts of responses to the possibility of unconscious/inaccessible meta-beliefs would require going well beyond Bondy & Carter's original objection.

Understood as a *sufficient* condition the doxastic view claims that: if I believe that *R* is my reason for believing that *p*, then my belief that *p* is based on *R*. (If you sincerely take *R* to be your reason for belief then it *is* your reason for belief.) So if as a result of the demonic appearances of normal basing, you believe that some normal, good reason *R2* is your reason to believe that *p* then it *is* your reason to believe that *p*. And so, according to Bondy and Carter, the intervention and mischief of the demon involving the bad reason *R1* would become irrelevant as the good reason *R2* would have become your reason to believe that *p*. But why should the bad reason *R1* become irrelevant? Bondy and Carter claim:

“if *S* believes *p* on the basis of both a good reason *R2* and a bad reason *R1*, *S*'s belief is justified by *R2*, and can count as an item of knowledge. So to be a sceptical scenario, the demon must make *S* believe *p* only on the basis of *R1*.” (Bondy & Carter, p.9)

Bondy & Carter here seem to take it as obvious that a belief that is simultaneously based both on one bad, non-justifying reason and one good, justifying reason can count as knowledge – i.e. the belief can be doxastically justified despite the presence of the bad reason that is also involved in the basing relation. But this is not at all obvious! Moreover, the possibility of a belief being based on multiple different reasons strikes me as interesting and worthy of discussion. Consider the following extreme (and so somewhat unrealistic) case: a subject, *S*, simultaneously bases her belief that *p* on 100 reasons, 99 of them bad and just one of them good. In being simultaneously based on 100 reasons, we can imagine that *S* ‘takes’ every single one of these 100 reasons to be part of her reason for belief. She is disposed to cite them all in support or defence of her belief, etc. For example, suppose *S* takes herself to have 100 independent mathematical proofs that *p*, each of which *S* takes to be individually a good enough reason for believing that *p*, and all of which *S* takes collectively to be the basis for her belief that *p*. But in fact 99 of those proofs are garbage and only 1 is genuine. I suggest that we feel the intuition here that *S* would not count as knowing that *p*. It seems that such a belief would not be doxastically justified and so not knowledge. All those bad reasons playing an

active basing role would seem to, as it were, overpower the one good reason and so infect the subject's basis for belief. Or suppose that a detective is equally disposed to cite 100 different reasons in support of her belief that Smith is guilty, 99 of which are crazy and bad, only 1 of which is a good reason. We surely have some intuition that the detective does not *know* that Smith is guilty, even if the one good reason on its own would have sufficed for knowledge. If that intuition is correct then Bondy & Carter are wrong to suggest that in order "to be a skeptical scenario the demon must make *S* believe *p* only on the basis of *R1*" and not also good reason *R2*. For once we allow beliefs to be based on multiple different reasons, it seems that a debasing Demon could make you believe *p* on the basis of thousands of different bad reasons at *t1*. Then if at *t2* one good reason (*R2*) also becomes a partial basis for the belief, so long as the thousands of bad reasons are also still partial bases for the belief, this belief would not count as 'an item of knowledge' and so we would have a skeptical scenario.

I think that whether or not one feels this intuition points us towards two importantly different notions of doxastic justification. According to the first notion: being doxastically justified in a belief is something like having used a valid permit or license for the belief. According to a second, different notion: being doxastically justified in a belief also requires not having used an invalid permit or license for the belief. Here is an analogy: suppose you have one genuine, valid fishing license and 99 invalid or fake fishing licenses. And suppose that at the entrance to the fishing lake you present all 100 of these 'licenses' to the fisheries official. One possibility is that all the invalid or fake licenses are just irrelevant so long as you presented the one genuine permit. But an alternate possibility is that presenting all the fake licenses to the official is an offence that gets you barred from fishing even though you also presented a genuine, valid license<sup>26</sup>. Here is one more analogy: the first notion of justification is kind of like using an old fashioned keyring with 99 useless keys and 1 key that opens the lock. So long as you used that 1 key you will have opened the lock and what you did with the 99 other keys is simply irrelevant. Whereas the 2<sup>nd</sup> notion of justification is more like a modern computer password system: if you inputted too many incorrect/fake passwords then the account is frozen and you are prevented from gaining access even if you also input the correct password. Here the use of fake keys matters to whether the use of a genuine key provides access.

I have suggested that there are at least some cases where our intuitions seem to go with the second notion of justification. Now, perhaps this sort of intuition could be resisted or rejected, but at the very least further argument is required – the first notion of justification (having used at least one valid permit) is not just obvious. The issues here thus turn on a third potentially controversial principle:

- **Potentially Controversial 3:** When a subject's belief is simultaneously based on many different reasons, if the proportion of Bad (non-justifying) to Good (propositionally justifying) reasons is sufficiently high then the subject's belief is not doxastically justified.

---

<sup>26</sup> Or consider licences to legally practice medicine. I suspect that in many jurisdictions, if you submit a bunch of fake, fraudulent medical degrees in your application to work as a doctor at a hospital this will get you barred from practicing medicine even if you also submit your one genuine medical degree as part of the application.

## 5. How not to worry about a so-called ‘Debasing Demon’

As mentioned, I do think that there is a genuine worry concerning the hypothesized demon’s systematic manipulation of the subject’s mental processes. Presumably the demon’s intervention in the basing process is supposed to make *no outwardly or introspectively detectable difference* whatsoever to the subject’s dispositions to behave, judge, assert, infer etc<sup>27</sup>. After all, the point of raising the scenario is that the choice between the hypothesis that we are basing our beliefs normally and the hypothesis that we are being debased is supposed to be *under-determined* by the evidence. The problem here, as I will now argue, is that such an undetectable intervention would not after all be debasing the subject but rather just adding an extra irrelevant casual epicycle into the belief-basing process.

If we think through how a debasing demon would operate, the demon would presumably somehow monitor the subject’s mental processes and intentions so that just as they are about to form a belief based in the normal way on the good propositional justifier, the demon intervenes and ensures that some other, bad, non-justifying reason is the basis for the belief, but then also has to create the false appearance that the belief is normally based and also erase any introspectable trace of the role of the bad basis. Since by hypothesis the subject would not notice any difference and would act, talk, behave, etc. exactly the same in the debasing scenario as in a normal, non-debased scenario, this intervention by the demon must have no outwardly observable effects nor any introspectively noticeable effects. But so then the worry is that the intervention of the so-called debasing demon and the involvement of the bad reason seems to be a difference that makes no difference. The subject would still publicly justify their belief to others by appeal to the good, propositionally-justifying reason. The subject would not be disposed to ever publicly cite the bad reason, or consider any hypothetical inference based on the bad reason, etc. So it seems reasonable to ask: in what sense is the subject’s belief really de-based? Why not think instead that the demon is just introducing an extra few bizarre, but epistemically irrelevant causal links in the chain between the propositional justifier and the belief? This extra causal wrinkle in the basing process happens to include a demon and a bad reason, but this is just a different, unusual way of *realizing* the proper basing relation that holds between the good reason and the belief. So rather than the good reason directly causing the belief, now the good reason causes a demon to spring into action, who somehow gets a bad reason to be causally involved in the formation of the belief, but who then also arranges things so that this intervention and involvement of the bad reason *has no further downstream traces in consciousness, memory, action, etc.*

As a comparison here consider yet another kind of possible demon, what one might be tempted to call a “De-Belief-ing demon”. A skeptic might propose a scenario as follows: ‘this new type of demon messes about with your mental processes so that although it still misleadingly *seems* that you have formed your beliefs as normal in response to your evidence and experiences, in fact you have not formed a belief at all! There is just the misleading introspective *appearance* of belief. The demon also produces the misleading appearance that it is your beliefs that are causing your speech and actions, your intentions and decisions and so on, when in fact it is the demon who is making sure that you continue to act and behave normally despite the lack of belief.’ Notice again that this skeptical scenario is meant to be perfectly compatible with all one’s evidence, so that the (alleged) “de-belief-ing” inflicted

---

<sup>27</sup> Conee (2015) uses the label ‘pure debasing’ for debasing that is not introspectively discernible.

upon the subject makes no outwardly or introspectably detectable difference. Insofar as we can make sense of being ‘de-belief-ed’ in this way, it would be a skeptical scenario. After all, a subject with no beliefs has no knowledge. However, I think the same sort of worry arises as to whether the demon’s intervention is really preventing beliefs from being formed, rather than just being a very unusual new way of *realising* the subject’s beliefs. Again, if everything seems exactly the same (experientially, introspectively), and you continue to behave in exactly the same way (verbally and non-verbally), it becomes unclear that the demon really has ‘de-belief-ed’ you. Why not say instead that the demon has now taken over the job from the parts of the brain that would normally play the role of realizing your belief? If we drew a functionalist ‘boxological’ diagram of the subject’s mental-causal system it would still look exactly the same, but now it would happen to be a demon that is physically realizing the belief box instead of the usual neural processes in the brain. So this so-called ‘de-belief-ing demon’ would actually just be functioning as a prosthetic or artificial belief.

Likewise, I suggest, the so-called ‘debasement’ demon would effectively be functioning as a prosthetic or artificial basing mechanism that is functionally and phenomenally identical to the subject’s normal, basing mechanism. The only differences that the so-called ‘debasement’ demon would make would be at the sub-personal level. But the involvement of the demon and the bad reason would make no difference to the personal level, functionally defined mental life of the subject. So whereas Bondy & Carter claimed that the deviant nature of the demon’s intervention would mean that the subject’s belief is not really based on the bad reason – which would simply leave the subject’s belief as based on nothing at all, so still debased – I have suggested that the demon’s intervention would in fact amount to a very bizarre and abnormal way for the subject’s belief to be based on the normal, good reason. And it is surely immensely plausible that the usual neurological processes, which happen to realize the basing relation in normal humans, are inessential to the basing relation – see McCain (2012).

At this point it is worth asking: is it an essential part of a debasement scenario that some external force, such as a demon, takes control of the subject’s belief forming mechanisms? The answer, I suggest, is: no. We can plausibly formulate debasement scenarios that do not rely on this. Firstly, we can perhaps coherently imagine a different kind of debasement demon, one that we might call a *Seducing Debaser*. This kind of debasement demon does not interfere with and take control over the internal mechanisms of the subject’s basing and belief-forming mechanisms. Rather, this demon *persuades* the subject to form a belief based on a bad reason rather than on the normal good reason, but then makes the subject immediately forget about this persuasion afterwards and creates the misleading introspective appearance that the belief was based on the normal, good reason. For example, perhaps the demon whispers such incredibly persuasive (though in fact specious) arguments for why the subject should form a belief based on the bad reason that the subject is convinced and forms the improperly based belief ‘on her own steam’, as it were. If this is something we can coherently imagine then it might allow for a debasement demon scenario in which the subject is still rationally responsible for the debased, defective nature of the belief<sup>28</sup>. Secondly, and more importantly, we need not posit any demon or external force at all. The demon is just a dramatic device. We could alternatively formulate a debasement scenario that only mentions the subject in question. The

---

<sup>28</sup> This Seducing Debaser would still be interfering with the subject’s mental processes insofar as it creates false memories and internal appearances of the belief’s basis. But a Seducing Debaser would not be taking over control of the belief forming process itself.



debasement skeptic could simply suggest: Consider your belief that *p* – perhaps you formed this belief based on a bad reason *R1* (or on no reason at all) but now your memory and introspection are playing tricks on you and it falsely seems to you that your belief is based on good reason *R2*. This scenario does not involve any kind of external manipulation of the subject's mental mechanisms – the skeptical hypothesis just appeals to the possibility that the subject herself formed a debased belief plus the possibility that her memory of and introspective access to the basis for her own belief is illusory and misleading.

What this shows, I think, is that the core worry about whether any debasing has even taken place does not concern the fact that the demon manipulates or takes control of the subject's belief forming mechanisms. Rather, the core worry is that – demon or no demon – the skeptical scenario requires that the supposed 'debasing' has no further introspectively or outwardly detectable effects. This is required insofar as the choice between the skeptical, debased hypothesis and the non-skeptical, non-debased hypothesis is supposed to be underdetermined by the subject's evidence. But it is precisely this lack of phenomenally accessible or functional/behavioral consequences that make it so implausible that any debasing would really have occurred. Even granting that a bad reason *R1* played *some kind* of causal role at time *t1* in the formation of the subject's belief, if by time *t2* the subject ends up being functionally and phenomenally identical to how she would have been anyway had she simply formed the belief based on good reason *R2* in the normal way, then it seems extremely plausible to maintain that at *t2* the subject's belief is not based on *R1* but rather on *R2*. Whatever the basing relation's precise nature, being based on one thing rather than another has to have some functional consequences. And if the causal involvement of *R1* in the belief forming process does not produce any of these functional consequences, consequences that could serve to differentiate a debased subject from a normal subject who based her belief on *R2* in the usual way, then *R1*'s causal involvement cannot count as *basing*.

On the other hand, if it is allowed that the subject in the debasing scenario is *not* functionally and phenomenally identical to the normal good basing case, well then the choice between the debased hypothesis and the non-debased hypothesis need not be under-determined by evidence and so we can reasonably expect to decide between these hypotheses based on our evidence in the normal way. Now we should expect a debased subject to behave differently from a non-debased subject in some way that reveals or evinces that the belief is based on *R1* rather than *R2*. In which case one can reasonably respond to the debasing skeptic as follows: I have not observed any of the distinctive consequences that the debasing hypothesis predicts. Rather, everything I have observed is as predicted by the normal, non-debased hypothesis. So the skeptical, debasing hypothesis is not something I need to worry about. This hypothetical debasing involving *R1* is supposed to make some potentially noticeable difference to my mental life or behavior concerning my belief that *p*. But as far as I can tell I have absolutely no inclination or disposition to rely on *R1* in justifying or defending my belief that *p*, nor have I ever yet behaved in any way that would suggest *R1* is the basis for my belief that *p*, nor do I currently believe that *R1* is a good reason to believe that *p*, on the contrary I think it is a bad reason to believe *p*, etc. And as far as I can tell all my inclinations and dispositions are to rely on *R2* in justifying and defending my belief that *p*, all my behavior suggests *R2* is the basis for my belief that *p*, I explicitly believe that *R2* is a good reason to believe that *p*, etc. Since the skeptical hypothesis makes different predictions to the non-skeptical hypothesis and since I have not observed any of these distinctively skeptical predictions, whereas everything I have observed is as predicted by the non-skeptical hypothesis, my total evidence straightforwardly

gives me very strong reason to prefer the normal non-debased hypothesis to the skeptical debased hypothesis. In other words, if we are supposed to imagine that the debasing would create behavioural or introspect-able effects that would be evidence for the debased hypothesis and against the normal, non-debased hypothesis, it seems we can reasonably respond that we simply have no such evidence. Rather, all our introspective and behavioral evidence seems to support the normal scenario. And so though it remains *possible* that our evidence may later turn out to be misleading, we have every reason to favour the normal non-skeptical hypothesis over the skeptical hypothesis.

In summary: it initially seemed like all that the debasing skeptic required is the assumption that knowledge requires proper basing and the assumption that this basing could be defective in some systematically unnoticeable way that would fit with all our evidence. But in this section I have argued that if the actions of the demon really are systematically unnoticeable in this way, then they must make no functional difference to the subject and so it is highly dubious whether any debasing has even taken place. On the other hand, if it is allowed that the actions of the demon do make some noticeable, functional difference to the subject, then the choice between the skeptical debased hypothesis and the non-skeptical, non-debased hypothesis is no longer one that is under-determined by our evidence. And so, assuming that no such noticeable sign of debasement has been observed, the non-skeptical hypothesis can reasonably be preferred over the debased hypothesis simply by appeal to what our evidence supports<sup>29</sup>. If the argument of this section is correct, it offers a more decisive response to debasing skepticism than that offered by Conee (2015). Conee accepts that an undetectable (or ‘pure’) debasing scenario is a genuinely skeptical possibility that would be experientially indistinguishable for the subject from a normal, non-debased scenario. He then goes on to argue that considerations such as economy, simplicity and coherence will still rationally favour the normal non-skeptical hypothesis over the debased skeptical hypothesis. Whilst I am myself generally in sympathy with fallibilist responses to skepticism, in the case of debasing skepticism I think that Conee’s line of thought is too concessive since, as I have argued, the undetectable ‘debasing’ demon scenario would not even be a genuinely skeptical debasing scenario in the first place. Assuming I am correct about this, it offers a stronger response to the debasing skeptic since it identifies a problem internal to the very idea of undetectable ‘debasing’ and does not require taking any stance on potentially controversial fallibilist theses about knowledge or justification.

## 6. Conclusion

I hope that the preceding discussion has clarified the skeptical threat that debasing scenarios pose and how the debasing-skeptic’s conclusion can be resisted. I suspect that most people who encounter the debasing demon scenario intuitively find it less worrying or gripping than

---

<sup>29</sup> Notice: assuming this response to the debasing demon is on the right lines it also offers a response to the alternative formulations of a debasing skeptical argument – discussed towards the end of section 2, above – which relied on the potentially controversial assumption PC1. For even granting PC1, we now have a basis to reject the first premise of these arguments. Since, to repeat, if the supposed debasing would create behavioural or introspect-able effects then, given the absence of any such effects, we *do* have a rational basis to prefer the normal, non-debased hypothesis to the Debasing Demon hypothesis. But if the supposed ‘debasing’ would not create any such behavioural or introspectable effects then we can just reject the idea that any debasing would have taken place even on the so-called ‘debasing’ hypothesis.

the traditional Cartesian scenario. The argument presented above in section 5 offered a new and improved reason to be unworried by the debasing demon scenario. To repeat: if the demon's intervention makes no functional or phenomenal difference then it becomes highly dubious whether any debasing would even take place. But if the demon's intervention is supposed to make some noticeable difference to either outward behavior or introspectable experience, then – so long as no such difference has in fact been observed – we can reasonably prefer the non-skeptical, non-debased hypothesis to the skeptical, debased hypothesis simply on the basis of what our evidence indicates.

Having presented my own preferred response to debasing skepticism, let me finish by briefly mentioning one further reason why one might feel that the debasing scenario is less worrying than the traditional Cartesian form of skeptical scenario. Let's now set aside the response to debasing skepticism presented above and suppose for the sake of argument that debasing really would take place in a debasing demon scenario. Even though (as we are now supposing) we would lack knowledge, we would still be *in touch with* external reality. We would still have (mostly) true beliefs and we would still have (mostly) accurate perceptual experiences. We are just supposed to lack this specific further positive epistemic property of belief – doxastic justification. (Moreover, the lack of this epistemic property would, by hypothesis, be totally undetectable and would have no further effects on behavior or consciousness whatsoever.) What is so troubling, I suggest, about the traditional Cartesian scenario is not just that we would lack knowledge, or even all the false beliefs we would have. It is that we would have *lost contact* with external reality and become 'trapped in our own heads'. This is, I suspect, a large part of why the Cartesian style of scenario, based on deceptive experiences, is so compelling: it vividly dramatizes the idea that we constantly live behind a 'veil of appearances'. The prospect of totally lacking knowledge, but still retaining all one's true beliefs and perceptual access to the external world is, I submit, much less worrying than the prospect of being trapped inside a bubble of false 'inner' appearances. Moreover, even if there is no Cartesian demon and our experiences are reliably caused by the external environment in the normal way, it can remain a troubling thought that the external world as it is in itself must remain forever *noumenal* and that our only access to it is indirect, via our own subjective appearances. I am not here endorsing or denying this familiar line of thought, but I do think that it is a large part of the reason why the classic Cartesian style skeptical scenario, much more so than the Debasing Scenario, exerts such a compelling grip on our imaginations<sup>30</sup>.

---

<sup>30</sup> Ancestral versions of this paper were presented at events in Pamplona, Dortmund, Liege, Oxford and Luxembourg. I am very grateful to the audiences on all those occasions for helpful feedback. I am especially grateful to Paul Silva for providing insightful written comments which very significantly improved the paper. My thanks also to Dietmar Heidemann, Guido Tana, Eva Schmidt, Simon Wimmer, Giulia Martina, Arnaud Dewalque, Matt Duncan, Frank Hofmann, Yannick Kohl, Anil Gomes, Mike Martin, Will Davies, Claire Field, Robin McKenna and the anonymous referees for this journal. Finally, I owe a special debt of thanks to Jonathan Schaffer, not only for writing the original paper on debasing skepticism, but also for very kindly and helpfully discussing an early version of this paper with me. Work on this paper was generously supported by the Luxembourg National Research Foundation (FNR) as part of the project 'The Epistemology of AI Systems', ref: C22/SC/17111440/EAI.

## References

- Ballantyne, N. & Evans, I. (2013) ‘Schaffer’s Demon’, *Pacific Philosophical Quarterly* 94 (4):552-559 (2013)
- Bondy, P. & Carter, J. A. (2018) ‘The basing relation and the impossibility of the debasing demon’, *American Philosophical Quarterly* 55(3): 203-215
- Brueckner, A. (2011) Debasing scepticism, *Analysis* 71(2): 295-297
- Carter, J. A. and Bondy, P. (eds.) (2020) *Well-Founded Belief: New Essays on the Epistemic Basing Relation*, New York, NY: Routledge
- Chalmers, D. (2022) *Reality+*, Penguin Books UK.
- Christensen, D. (2010) Higher Order Evidence, *Philosophy and Phenomenological Research* 81 (1): 185-215
- Conee, E. (2015) ‘Debasing Skepticism Refuted’, *Episteme* 12(1): 1-11
- Cunningham, J. (2021) ‘The Basis of Debasing Scepticism’, *Erkenntnis* 86: 813-833
- Das, N. & Salow, B. (2018) ‘Transparency and the KK principle’, *Noûs* 52 (1):3-23
- Descartes, R. (1984) *The Philosophical Writings of Descartes*, (eds.) Cottingham, J., Stoothoff, R. and Murdoch, D., Cambridge University Press.
- Dorst, K. (2019) ‘Abominable KK Failures’, *Mind* 128 (512): 1227-1259
- Dretske, F. (1970) ‘Epistemic Operators’, *Journal of Philosophy*, 67: 1007-23
- Feldman, R. & Conee, E. (2007) ‘Making Sense of Skepticism’, in Conee E. and Feldman R. (eds), *Evidentialism: Essays in Epistemology*, Oxford University Press.
- Goodman, J. & Salow, B. (2018) ‘Taking a Chance on KK’, *Philosophical Studies* 175 (1):183-196
- Greco, D. (2014) ‘Could KK Be OK?’ *Journal of Philosophy* 111 (4):169-197
- Hawthorne, J. (2005) ‘The Case for Closure’, in *Contemporary Debates in Epistemology*, Steup, M. & Sosa, E. (eds.), Blackwell Publishing.
- Horowitz, S. (2014) ‘Epistemic Akrasia’, *Noûs* 48 (4): 718-744
- Korcz, K. (1997) “Recent Work on the Basing Relation”, *American Philosophical Quarterly*, 34(2): 171–191
- Lasonen-Aarnio, M. (2014) ‘Higher-Order Evidence and the Limits of Defeat’, *Philosophy and Phenomenological Research* 88 (2): 314-345.
- Lasonen-Aarnio, M. (2020) ‘Enkrasia or evidentialism? Learning to love mismatch’, *Philosophical Studies* 177 (3): 597-632
- McCain, K. (2012) “The Interventionist Account of Causation and the Basing Relation”, *Philosophical Studies*, 159: 357–382
- McCain, K. (2020) “What the Debasing Demon Teaches Us About Wisdom”, *Acta Analytica* 35: 521–530.
- Nelson, M. (2010) ‘We Have No Positive Epistemic Duties’, *Mind* 119 (473): 83-102
- Neta, R. (2019) “The Basing Relation,” *Philosophical Review*, 128(2): 179–217
- Nozick, R. (1981) *Philosophical Explanations*, Harvard University Press.
- Pryor, J. (2018) ‘The Merits of Incoherence’, *Analytic Philosophy* 59 (1): 112-141
- Putnam, H. (1981) *Reason, Truth and History*, Cambridge University Press.
- Russell, B. (1912) *The Problems of Philosophy* Oxford University Press.
- Schaffer, J. (2010) ‘The Debasing Demon’, *Analysis* 70(2): 228-237
- Silva Jr., P. (2015) “Does Doxastic Justification Have a Basing Requirement?”, *Australasian Journal of Philosophy*, 93(2): 371–387
- Silva Jr., P. (2019) ‘Beliefless Knowing’, *Pacific Philosophical Quarterly* 100 (3): 723-746.
- Silva Jr., P. (2022) ‘Does the Basing Demand on Doxastic Justification Have Dialectical Force? A Response to Oliveira’, in Oliveira & Silva (eds.) *Propositional and Doxastic Justification: New Perspectives in Epistemology*, Routledge.
- Smithies, D. (2019) *The Epistemic Role of Consciousness*, Oxford University Press.
- Sylvan, K. (2016) “Epistemic Reasons II: Basing”, *Philosophy Compass*, 11(7): 377–389
- Tal, E. (2020) ‘Self-Intimation, Infallibility, and Higher-Order Evidence’, *Erkenntnis* 85(3): 665-672
- Titelbaum, M. (2015) ‘Rationality’s Fixed Point’ *Oxford Studies in Epistemology* 5
- Vogel, J. (1990) ‘Cartesian Skepticism and Inference to the Best Explanation’, *Journal of Philosophy* 87(11): 658-666.
- Weatherson, B. (2019) *Normative Externalism*, Oxford University Press.
- Wedgwood, R. (2022) ‘Doxastic Rationality’, in Oliveira & Silva (eds.), *Propositional and Doxastic Justification: New Essays on Their Nature and Significance*, Routledge.
- Williamson, T. (2000) *Knowledge and its Limits*, Oxford University Press