Engineered knowledge, fragility and virtue epistemology

Daniel O'Brien

There is a clean image of knowledge transmission between thinkers that involves sincere and reliable speakers, and hearers who carefully assess the epistemic credentials of the testimony that they hear. There is, however, a murkier side to testimonial exchange where deception and lies hold sway. Such mendacity leads to sceptical worries and to discussion of epistemic vice. Here, though, I explore cases where deceit and lies are involved in knowledge transmission. This may sound surprising or even incoherent since lying usually involves saying something that is false. Even if a liar unwittingly tells the truth, those believing their testimony would not acquire knowledge. In §1, however, I suggest a range of examples where lies can be used to transmit knowledge. The kinds of cases upon which I focus are those involving 'engineered knowledge', where a speaker skilfully manipulates another into having a true belief. Such cases are epistemologically significant for several reasons. First, they are counter-examples to a range of widely held assumptions about knowledge. We have, for example, cases where sincerity is not necessary for knowledge transmission, where there is knowledge from falsehood, and, as we shall discuss in §2, knowledge that is fragile or not safe. Second, in §3 and §4 I suggest that engineered knowledge is best captured by a virtue theoretic approach to testimonial transmission.

1 Engineered knowledge

Iago. In Shakespeare's play, Iago lies to Othello that he has seen his rival, Cassio, wiping his face with a fancy handkerchief in order that Othello comes to believe that *Desdemona is not in possession of her handkerchief*.[1] Iago has seen no such thing—in fact, he has the handkerchief,

[1] I have played a little fast and loose here with the plot. Iago actually says that he has seen Cassio with a handkerchief that looks like Desdemona's, rather than merely a fancy handkerchief. It could be inferred from the original dialogue that Desdemona's handkerchief is missing, given certain other plausible premises (perhaps the handkerchief is of an unusual kind and there are likely to be very few like it in Venice). Shakespeare's Iago also lies, but my

after taking it from Desdemona's maid, Emilia. Othello's acquisition of this true belief about Desdemona not being in possession of her handkerchief is ultimately in the service of deception and the lie that Desdemona is having an affair with Cassio, but it is a true belief nonetheless and so is at least a candidate for knowledge. Iago has not seen Cassio with the handkerchief and so his utterance is a lie. Iago surmises that Othello will come to believe that Cassio has Desdemona's handkerchief and thus that Desdemona is not in possession of it, given Othello's jealousy of Cassio and the significance of the handkerchief (it was Othello's first gift to his wife).

Cases of engineered knowledge typically have three properties. First, there must be a suitably motivated and devious speaker. Second, the hearer must have a cognitive failing of some kind. Othello's failing is pathological jealousy that leads to gullibility and credulousness. Almost anything would suggest to Othello that his wife is unfaithful.

> …Trifles light as air
> Are to the jealous confirmations strong
> As proofs of holy writ. (Act 3, Scene 3)[2]

Third, it must be the case that the hearer's beliefs do not turn out true by accident. It is Iago's intention that Othello comes to have this belief about the handkerchief and through his nefarious activities and devious testimony he is able to transmit this information to Othello.

This is not a Gettier case, one in which a thinker has a justified true belief by accident. Turri (2012, p. 248) characterises Gettier cases as involving 'double luck' (see also Zagzebski, 1994, p. 66): 'Start with a belief sufficiently justified to meet the justification requirement for knowledge. Then add an element of bad luck that would normally prevent the justified belief from being true. Lastly add a dose of good luck that "cancels out the bad," so the belief ends

---

example gives more emphasis to the role played by Othello's irrational jealousy: Iago's utterance leads to Othello believing that his wife's handkerchief is missing even though there's no good reason to think this. See [Anonymised Reference] for extended discussion of this example.

[2] Jealousy is an epistemologically interesting emotion: on the one hand, it may lead a person to have detailed knowledge of their partner's movements and acquaintances, but, as in my example, it is accompanied by the epistemic vice of credulity.

up true anyhow.' This double dose of luck precludes attribution of knowledge in such cases. Testimony concerning the whereabouts of a handkerchief would usually provide justification for beliefs concerning whether or not someone is in possession of such an item, but in this case it could be thought that Othello is prey to bad luck since Iago is lying. This bad luck is cancelled out by Iago's intentional strategy to make it the case that Othello acquires the belief that Desdemona in not in possession of the handkerchief. It is not, though, appropriate to describe such a case as involving good and bad luck in this way. Othello is told a lie *in order* that he comes to acquire a true belief about the whereabouts of the handkerchief. This was Iago's intention *from the start*, and the success of his actions is down to his cunning and intelligence and not to luck.[3]

Violin. Parents lie to their children, and often for good reason. Such lies are termed paternalistic. Gerald Dworkin (2014) discusses a certain kind of these, ones that are an 'efficient means of producing a kind of experience which . . . leads the deceived person to realize he has certain abilities and capacities he did not recognize he possessed.' These are lies that 'involve getting someone to do something he already has the ability to do. He just does not know it.' This he calls The Dumbo Effect. In the Disney film, Dumbo has the ability to fly yet does not believe he can. Mouse gives Dumbo a magic feather, lying to him that this can make him fly. This seems to work and Dumbo flies. Later on he drops the feather and, thinking that it's only this

---

[3] Gettier cases can be embedded in engineered scenarios, with the result that the crucial element of luck is eliminated and there is therefore no reason to deny that knowledge is acquired in such embedded cases. Consider the following twist to one of Gettier's original examples. The boss, for some reason, wants Smith to know that the man who will get the job has ten coins in his pocket. To do this he plants ten coins in the pockets of both Smith and Jones, tells Smith that Jones will get the job, and engineers it so that Smith comes to know the contents of Jones' pocket. I suggest that the correct intuition here is different to that in the original Gettier case. As originally described, Smith was exposed to bad luck in that the boss lied to him about Jones being in line for the job. However, this bad luck was cancelled out by Smith's good luck, that being, that he also had ten coins in his pocket and the fact that he was given the job. In my embedded Gettier case, it is not bad luck that the boss lies to Smith. The lie is part of the boss's strategy to transmit a true belief to Smith, the true belief that the man who will get the job has ten coins in his pocket.

that enables him to fly, he falls out of the sky. Mouse then reveals that he had lied to Dumbo about the feather and this then enables him to take flight once more, realizing that this is something he has the ability to do. This is one kind of lie parents tell to their children. Your violin practice always sounds great . . . You're much more grounded than I was at your age . . . Lies—but ones that help a child realize, and, I claim, know, that they are good enough to get into the school orchestra and that there's no reason to be worried about their lack of direction.

Attic. Another kind of paternalistic lie can play a protective role. These can involve a testifier citing evidence they know to be false or using an argument they know to be invalid. The attic is a dangerous place for a young child and this is something that needs to be taught. One way to do this is to say that monsters live up there.[4, 5]

In *Iago*, *Violin* and *Attic* the speaker acquires true beliefs via testimony and this is no accident; the speaker engineers the situation so this is so. It is not the case, though, that non-accidentally acquired true beliefs always amount to knowledge, even if, in a case of testimony, it is the intention of the speaker that the hearer comes to acquire beliefs that the speaker knows to be true. More therefore needs to be done to substantiate the claim that knowledge is acquired

---

4 Implausible, rather precious, claims are sometimes made concerning the morality of lying. Bok (1999, p. 31), following Kant, claims that '[o]nly where a lie is a *last resort* can one even begin to consider whether or not it is morally justified.' But, as Nyberg (1993, p. 26) puts it: 'Haven't we got the value of truth telling just a little out of focus?' It's easy to think of cases where it is right for a parent to lie to their children about, say, divorce or the precariousness of the family's finances, and this is so even if that's not the last resort. Dworkin (2014) points out that this claim is nicely made by the novelist Michael Chabon (2013): 'When a young child's fingers brush against a crack in the world, a parent seeking to account for that fragility may lie or tell the truth. Either is permissible, depending on the circumstances.'

5 The cases of engineered knowledge discussed involve knowledge from falsehood. As such, they are not consistent with counter-closure, that is, the principle that: 'If (i) *S* knows *q*, and (ii) *S* believes *q* solely on the basis of a competent deduction from some premises including *p*, then (iii) *S* knows *p*' (Ball & Blome-Tillmann, 2014, p. 552). For others who argue for knowledge from falsehood and against counter-closure, see Warfield (2005), Klein (2008) and Fitelson (2010).

in these cases. A key claim is that these are cases of *engineered* knowledge and they therefore require careful planning and execution. Such engineering is therefore—in some sense—a cognitive achievement on the part of the speaker and this claim will be crucial in §3 and §4 where I suggest that such cases are usefully seen through the lense of virtue epistemology.

One response to these cases of engineered knowledge could be to attempt to account for them by appealing to the distinction between semantics and pragmatics, between the stable meanings of words and what we do with them. Iago's claim concerning the fancy handkerchief has a literal meaning, and it has further pragmatic force when uttered to Othello; similarly with respect to the dangerous attic. Cases of engineered knowledge, however, cannot be subsumed under such an account, at least as construed by Grice (1989). Where there is Gricean implicature, the hearer picks up on the speaker's intention and thus comes to take the assertion in the way intended. In irony, for example, one communicates something one believes to be true by saying something that one believes to be false, and the hearer recognises this. If this were so in the cases I discuss, the deceptive element would drop out since both parties would have knowledge of the relevant implicatures. Othello, however, does not recognize Iago's intention, nor does the child curious about the attic, or the violinist. There is an art to both irony and engineered knowledge: irony involves signalling one's intentions in order to make it clear that one should not be taken literally, perhaps with a cough or a wink, or with vocal emphasis or intonation. Iago's art, in contrast, and that of the parent, involves concealing their strategy from the recipient.

For similar reasons engineered knowledge should not be construed in terms of Bach's (1994) impliciture or Carston's (2002) explicature. In both, the content of an utterance goes beyond what is said, but does not amount to implicature. In implicature the speaker means something by his literal utterance ('8 p.m. is an unusual time for dinner') and also intends to communicate something else ('that's too late'); in impliciture nothing beyond the meaning of the utterance is conveyed, but more is communicated than that covered by the logical form of the sentence spoken: 'Jack and Jill are married' is said, but the impliciture is that they are married to each other—and this is also meant. But in these cases too, the hearer assumes that the speaker intends for her to figure out the enrichment or expansion of what was said. Engineered knowledge, though, is not transparent: the hearer is not party to the speaker's intentions. Further, the communicated content is not objective in the way that it is in the case of implicature, impliciture and explicature. Engineered knowledge can be more specifically targeted, the speaker's statements engineered for the specific knowledge and epistemic weaknesses of the hearer. Fricker (2006, p. 599) is therefore wrong to claim that 'telling is like

radio broadcasting. Anyone who can tune her radio into the wavelength can pick up the signal'. This is not so with respect to engineered knowledge.

There are therefore a range of examples where lies and deception can transmit true beliefs.[6] Othello comes to believe the handkerchief is missing, the budding violinist that he is capable of getting into the orchestra and the child that the attic is dangerous. In what follows I will go on to substantiate the claim that it is not simply true belief that is transmitted, but knowledge.

2 Safety and fragility

Almost the only thing upon which epistemologists agree is that in order to have knowledge it cannot be the case that one's true beliefs are acquired by accident. A certain kind of luck must therefore be ruled out, that is, veritic luck: it must not be accidental that a belief turns out to be true given the evidence that one has for that belief.[7] The leading anti-luck theory concerns

---

6 Welbourne (2001, p. 102) claims that knowledge transmission 'requires that she [the speaker] is not making an improper use of the speech-act that she performs', arguing that knowledge cannot be transmitted by, for example, a double bluff: in such a case the 'speaker can hardly be described as sincere and I doubt anyone would think that either knowledge or even a good quality belief could be obtained *from* a source which is so contaminated by deceit'. Such deceit, though, could also be used by an engineer. A teenager justifiably believes that his parents will assume he's lying about who broke the vase at the party, and that they are also aware that he is likely to know they will take his utterance in this way. He says it was Craig (which it was). His parents would take him to be lying, but for the fact that they think he knows they will be sceptical of his claim, and so they come to believe it was Craig. This, I claim—contra Welbourne—is something that they now know, given that the teenager was one step ahead of his parents and aware of how his utterance would be taken.

7 See Pritchard (2005) for discussion of the various species of epistemic luck. There are other forms of luck that are not epistemically damaging: capacity epistemic luck, where 'it is lucky that the agent is capable of knowledge' (ibid., p. 134) and evidential epistemic luck, where 'it is lucky that the agent acquires the evidence that she has in favour of her belief' (ibid., p. 136);

safety—the basic idea being that one cannot know something that could have easily been false. This is usually cashed out in terms of what one would believe in nearby possible worlds. Pritchard (2009, p. 34) puts it thus:

> S's belief is safe iff in most near-by possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world, and in all very close near-by possible worlds in which S continues to form her belief about the target proposition in the same way as the actual world, her belief continues to be true.[8]

McCain (2014) argues that *Iago* does not satisfy this condition. Othello's belief that Desdemona's handkerchief is missing is not safe.[9] There are nearby worlds in which his belief is acquired in the same way, but where this belief is not true. Othello's belief could easily have been false. It only matters to Iago that the claim about the handkerchief is true in so far as this is useful to him. His strategy may be more precarious if the belief transmitted to Othello were false, since, if the handkerchief were not missing, Othello could happen upon it and there would be a danger that Iago's story might start to unravel. Iago, though, is cunning, and he would undoubtedly concoct further lies to explain why the handkerchief remains or seems to remain in Desdemona's possession. There are, then, nearby worlds where Othello's beliefs are formed in the same way as in the actual world—that is, on the basis of Iago's testimony—but where those acquired beliefs are false. McCain argues that such beliefs are therefore not safe and do not amount to knowledge.

Paternalistic lies also do not lead to safe beliefs since a child would believe almost anything a parent says whether or not the attic is dangerous and whether or not their violin playing is impressive. There are nearby worlds in which a child comes to believe he can get into the orchestra on the basis of his parents' praise even when this is not the case, and where

---

she may, for example, just have been there at the right place and time to see and come to know that a badger lives on the embankment.

8 For other accounts of safety see Sosa (1999) and Williamson (2000).

9 This belief is also not *sensitive* given that Othello would have the belief even if the handkerchief is still in Desdemona's possession. A belief is sensitive if it is the case that a thinker believes *p* when *p* and does not believe *p* when not-*p*. See Nozick (1981).

he believes the attic is dangerous even though it is not (there may be other reasons the child should not go there).

I shall consider two responses to McCain's objection. First, one could reject his claim and argue that in some cases engineered beliefs are safe. As said, counterfactuals constituting safety are cashed out in terms of what would be believed in nearby possible worlds. McCain, I think, assumes that nearby worlds must include those where Iago lies and where the handkerchief is not missing, and—although he does not discuss this example—where a parent lies and where the attic is in fact not dangerous. Across such worlds safety is compromised. There is, though, another way of thinking of the proximity of worlds in these cases. After all, the key feature of some of them is that a speaker is lying in order to inform a hearer of a certain truth. Nearby worlds could therefore be seen as those in which that truth still holds, but where other features of the world are variable: worlds, for example, where the attic is dangerous, but where children are brave and climb the attic-stairs regardless. In the case of paternalistic lies, a world in which a child gullibly believes a parent's purposeless, rather than engineered, lies seems to be farther away—in the space of (intuitive) possibilities—than one in which a parent continues to attempt to communicate (protective) truths to their children, but where other aspects of the situation vary. On this interpretation of nearby worlds, such engineered beliefs could be seen as safe since the worlds in which the relevant beliefs are false would not be nearby. (Such a response is not so plausible in the case of *Iago* since the engineering here is not performed in the service of truth. Iago, unlike the parents in *Violin* and *Attic*, is not primarily concerned with whether Othello has a true belief about the handkerchief; Othello's jealous response is all that matters.)

A second reply to McCain accepts that Othello's beliefs are not safe, but denies that this impugns their status as knowledge. We should remember here why the safety condition on knowledge is introduced. It is a condition designed to rule out veritic luck. Epistemic dogma is hard to resist, and it is fed by ingenious examples involving fake barns, shaggy dogs, benevolent demons and that panoply of residents of the world of contemporary analytic epistemology. Responses to more mundane cases can then be distorted by mantras such as 'No Knowledge from Falsehood', 'No Safety, no Knowledge'. . . . But let us approach engineered knowledge with an open mind; let us look and see. Surely, for example, there's nothing clearer than the child knows the attic is dangerous. Or, at the very least, one's commitment to the claim that the child has knowledge is stronger than one's commitment to the safety condition on knowledge. And why does the child have knowledge? Well, because of the engineer. Their intervention explains why these are not cases involving merely lucky true beliefs, even though

this explanation is not one that is 'modally robust'. Riggs (2012, p. 278) reminds us that 'luck matters because knowledge is an accomplishment': to say that a true belief is lucky is to say it is 'out of his or her control—it is not something that the agent brings about' (ibid., p. 292). Understanding luck in this way enables us to attribute knowledge in engineered cases because—even if the relevant beliefs are unsafe, insincere and derived from falsehood—their transmission is under the control *of the speaker.* Riggs delineates two kinds of luck that might endanger such control: where the outcome, in this case a true belief, is not causally due to the agent's abilities, and where the outcome is not intended. The role of the engineer rules out both kinds of luck. The crucial step here—one that Riggs himself does not take—is to see the relevant agent as the speaker and not the recipient of testimony, and therefore to embrace the deep social nature of testimonial exchange.

Safety is only one way of articulating the anti-luck condition on knowledge. Othello's belief is not acquired by accident. Iago believes that Othello's jealousy will be fired by his report and thus his awareness of this facet of Othello's psychology allows him, I claim, to transmit knowledge concerning the handkerchief. It is, in fact, such lack of safety upon which the speaker preys in such cases. The hearer's gullibility—the fact that he could easily have such beliefs even if they were not true—allows the speaker to engineer beliefs in the ways described. One should therefore reject Pritchard's claim that safety is a necessary condition on knowledge. Knowledge can be unsafe.[10]

The conception of knowledge with which I am working is the following: knowledge is non-accidental true belief that is the result of a cognitive achievement. This is a conception of knowledge shared by, for example, Sosa and Greco. 'knowledge is true belief out of intellectual virtue, belief that turns out right by reason of the virtue and not just by coincidence' (Sosa, 1991, p. 277); '*S* knows that *p if and only if S* believes the truth (with respect to p) because *S's* belief that *p* is produced by intellectual ability' (Greco, 2010, p. 71). My twist, though—in cases of engineered knowledge—is that the epistemically important abilities and virtues are

---

[10] Goldberg (2005) has argued that testimonial knowledge can be acquired from unsafe testimony. He presents an example intended to show that a testimonial belief can be safe, even if the testimony on which it is based is unsafe: the lack of safety need not transmit from the testimony to the testimonial belief. Note, though, that the suggestion in my case is distinct from this. It is not just that Iago's testimony is unsafe, but so is Othello's belief—yet I still want to claim this is a case of knowledge transmission.

those possessed by the engineer and not the subject. In §4 I provide a more detailed account of the role of the engineer.

Sosa (2007), who we will return to below, also allows for the possibility of unsafe knowledge; knowledge that is not modally robust:

> Knowledge is simply . . . apt performance in the way of belief. Knowledge . . . does not require the safety of the contained belief, since the belief can be unsafe owing to the fragility of the believer's competence or situation. (ibid., p. 41)[11]

Perceptual knowledge can be fragile in this way. This is the case, Sosa claims, given the Cartesian possibility that in nearby worlds we could be dreaming.[12] Engineered knowledge can also be fragile. There are nearby possible worlds in which Othello acquires his belief about the

---

11 Sporting analogies lend themselves well to discussion of such fragility. The skill of sportsmen and women can be fragile without this diminishing the craftsmanship of their successful performances. One thinks here of Alex Higgins' long pots and David Gower's square cuts—admittedly rather parochial British examples from my youth, the fragility of which made them all the more appealing.

12 Various forms of fragility make an appearance in the literature. Graham (2014) argues that perception is not reliable across all possible worlds and it may turn out that '[a] perceptual process may even be reliable in just one locale, without being reliable in most environments the creature might really traverse. Reliability in those other locales might not matter to the creature. That's why perception can be non-accidentally reliable without being transglobally reliable, let alone globally reliable' (ibid., p. 532). Comesana's *Halloween Party* (2005) involves coming to have knowledge of the whereabouts of a party even though there are nearby worlds in which one does not have the relevant true belief. Gaultier (2014) argues that beliefs can be safe in a good environment even if they are not safe in a bad environment. There may be a sense in which it is lucky that one happens to be in a good environment, but this is not epistemically pernicious luck: 'a true belief can at once constitute knowledge and be unsafe, but only in the sense that it could easily have been formed in a different environment in which it would have been false, and not in the sense that it could easily have been false in the environment in which it has been formed' (ibid., p. 488). Also see Hetherington (1999).

handkerchief in the same way but where this belief is false. There's a nearby world where the same lies are told to a child but where the attic is not dangerous.[13]

Engineered knowledge may be fragile, but the true beliefs acquired are not prey to epistemically corrosive forms of luck. The anti-luck condition is satisfied in virtue of there being, on particular occasions, an explanation of epistemic success that refers to the skills of the speaker. The beliefs of Othello and of the child are not true by accident, even though they are not safe. Iago believes that Othello's jealousy will be fired by his report and thus his awareness of Othello's mind allows him to transmit this knowledge concerning the handkerchief to Othello. Iago's scheming should be seen as eliminating veritic epistemic luck since it is not down to chance that Othello's belief is true—Iago engineered the situation so that this would be so. Similarly, the parent knows their child's likely reaction to talk of monsters.

3 Credit, hidden helpers and epistemic angels

Pritchard also argues for a second condition on knowledge—the ability condition—and consideration of this will support my suggestion that luck is ruled out in engineered cases by the epistemic activity and abilities of the speaker.

---

13 Others have suggested cases where knowledge is acquired via deception or lies or when the speaker asserts what she wrongly believes to be false. Audi (2013) and Klein (2008) discuss a case where a child comes to know that presents will arrive in the morning on his parents' testimony that Santa will bring them. Lackey (1999, pp. 47–59) discusses two cases, *creationist teacher* and *consistent liar*, in which false beliefs are involved in testimonial transmission and where insincerity is involved on the part of the speaker. Audi claims that the false belief about Santa is not essential for knowledge transmission and that the child's knowledge is based on 'parental assurance' (2013, p. 515) that there will be presents. This is reliably grounded true belief that is not testimony-based, but arises '*by way of* testimony, even when the attestation is false' (ibid., p. 516). Lackey also grounds the knowledge acquired in the reliability of the claims that the teacher makes. I shall argue below, though, that engineered knowledge should not be seen in reliabilist terms (discussed here in terms of safety), but rather in terms of certain epistemic virtues possessed by the engineer.

The ability condition is motivated by consideration of *Temp*. Through careful observation of a thermometer, Temp has reliable beliefs about the temperature. The thermometer, though, is broken and fluctuates randomly. There is, however, 'an agent hidden in the room' (Pritchard, 2012a, p. 260) who adjusts the thermostat so that the reading on the thermometer is correct. Temp's beliefs are therefore reliable and safe, but Pritchard argues that the intuition we should share is that they do not amount to knowledge since this reliability is not down to Temp; he cannot be *credited* with it. Pritchard claims that there must therefore be a further condition on knowledge: it must be 'the product of one's cognitive abilities, such that when one knows one's cognitive success is, in substantial part at least, creditable to one' (2010, p. 51).[14] He incorporates this into his 'anti-luck virtue epistemology': knowledge requiring both safety and the exercise of epistemic virtue.

Pritchard's ability condition may appear problematic for my examples given that it is not the case that those being engineered should be given credit for their true beliefs. I do, however, endorse the spirit of this condition—knowledge must be creditable to an agent—but this, I argue, need not be the knower themselves. Engineered knowledge cannot satisfy Pritchard's interpretation of this condition given the cognitive failings on the part of the knower. Knowledge is nevertheless acquired in such cases and this is due to certain skills of deception on the part of *the speaker*.

Pritchard argues that Temp does not have knowledge of the temperature because of the intervention of a 'hidden helper' (2012a) and thus his true beliefs are not to his credit. In an earlier paper he offers a different dramatization of the *Temp* scenario, one in which the alter-ego of Descartes' evil demon adjusts a thinker's beliefs so they are true, regardless of how irresponsibly they may be acquired. Again, though, knowledge should not be ascribed to such a thinker since their 'cognitive success in no way relates to . . . cognitive ability' (2006, p. 20).[15]

---

[14] Also see Greco (2003, p. 111): 'To say that someone knows is to say that his believing the truth can be credited to him. It is to say that the person got things right due to his own abilities, efforts and actions, rather than due to dumb luck, or blind chance, or something else.'

[15] Levy does not share Pritchard's intuition: 'If the intervention is a stable feature of Temp's environment—which it appears to be, as the case is described—then Temp forms his belief by consulting a thermometer in perfectly good order. Admittedly, the thermometer doesn't work in the way he might imagine it does, since he is ignorant that another agent manipulates it so that its readings match the actual temperature, but ignorance of the workings of the epistemic

Engineered knowledge also involves a kind of hidden helper: the agents themselves are not literally hidden from view—as in the cases of the demon and Temp—but they are hidden in the sense that their deceptive strategies are concealed. Pritchard would not therefore see my cases as involving knowledge. Here, though, I present four considerations that undermine the intuition that knowledge is not acquired in (at least some) cases where there are hidden helpers.

First, our intuitions may be affected by thinking of the helper (or engineer) as a 'demon' (albeit a benevolent one). On my reading, the agent envisaged is more akin to an epistemic angel. Bobier (2014) discusses the intervention, not of epistemic angels, but of God, and his intuitions concerning this case are contrary to those of Pritchard. Godfrey comes to believe in God after he appears to him in a dream. Bobier suggests that this is plausibly a case of knowledge since 'God, being all-powerful, can surely usurp an unreliable method in order to reveal in a knowledge-bestowing manner himself or his will to someone' (2014, p. 312). Godfrey, however, deserves little credit for his belief, and so Bobier takes this scenario as telling against Pritchard's anti-luck virtue epistemology since the ability condition is not satisfied. (Godfrey's belief may, however, satisfy Pritchard's safety condition since '[i]n every nearby possible world God exists, and the proposition <God exists> is true. Thus, in every nearby possible world where Godfrey . . . form[s] the belief <that God exists> on the basis of a dream, . . . [his] belief will be true, and thereby, safe' (ibid., p. 313). Assuming the necessary existence of God, safety is built into the scenario in that the belief that he exists cannot but be true in all nearby possible worlds.)[16]

---

devices we rely on is ubiquitous. I don't know how my smartphone works, but I seem to know a great many of the facts that I retrieve from it (for that matter, I have only the sketchiest understanding of how my brain works; that doesn't prevent me from gaining knowledge, for instance by retrieving memories)' (Levy, 2014, p. 10). Levy does not find Pritchard's intuition persuasive. He is, though, taking the relevant epistemic devices to be reliable and it is this, I think, that makes him confident that we are talking about knowledge in these cases. I have suggested, though, that knowledge can be engineered even in the absence of such reliability (I have discussed reliability in terms of safety).

[16] One may be reminded here of my first response to McCain above, where the beliefs that parents engineer are by their nature safe since the relevant nearby worlds are those in which such beliefs are true, given the overarching goals of parenthood.

There is, though, an alternative way to uphold the intuition that Godfrey has knowledge of God, and that is to focus on the epistemic abilities, not of the hearer (Godfrey), but of the 'speaker' (God). In Bobier's case knowledge acquisition is secured by the actions of God and in cases of engineered knowledge it is secured by the actions of the engineer. Bobier only focuses on Godfrey's particular belief concerning the existence of God, but God could usurp our unreliable methods more widely, thus leading to a scenario akin to that of Pritchard's benevolent demon. Approaching the latter scenario from this direction—working up from the claim that God could 'reveal in a knowledge-bestowing manner himself or his will to someone'—may enable us to see that knowledge can be ascribed even where there are hidden helpers and engineers. Bobier does not therefore need to reject virtue epistemology; it is, though, the virtues of the engineer that should be our concern in these kinds of cases, and in Bobier's case these are certain divine intellectual virtues.

Second, we need to be careful to distinguish two kinds of cases. First, there could be angels that bring the world in line with our beliefs. This is the version of the scenario that Pritchard describes: 'every time I form a belief the world is adjusted so that my belief is true' (2006, p. 20). Thus, 'I am not being responsive to the facts at all, but rather the facts are being responsive to me' (ibid.). Here, then, it may not seem appropriate to ascribe knowledge.[17] There is, however, a distinct kind of case where the world is thus and so, and where an angel engineers it that we have true beliefs about this mind-independent realm. With the help of an engineer our beliefs would thus track the world. It is not so clear that ascriptions of knowledge are out of place in such cases.[18]

---

17 Even in this case, though, I do not think it is completely obvious that knowledge cannot be attained. Lewis (1973, p. 84) is also somewhat ambivalent when discussing similar cases involving self-fulfilling beliefs—where, for example, a child's beliefs about what he will get for Christmas are made true by his parents even though they had not originally intended to give him those presents: 'The situation is a peculiar one, and my intuitions, and I would suppose other people's, are not completely clear on the matter. But it seems, on the whole, that we ought not to speak of knowledge here. The essential point of a "faculty of knowledge", is that it should, in respect of what is known, be passive to the world. If the "reflection" is achieved by our mind moulding the world, we are not knowing but creating.'

18 Iago's engineering runs deep. First, he does not merely work with the jealousy he finds in Othello; he is, rather, instrumental in creating this aspect of Othello's character. Thanks to

Third, let us consider early modern occasionalism. For Malebranche (1997), finite substances are causally impotent; only an infinite substance, or God, has causal efficacy. God therefore mediates all causal interactions: he makes it the case that the red ball moves away from the cue ball after it is struck. He also therefore mediates the causal interactions involved in knowledge acquisition: those between physical states of the brain, those involved in mind-body interaction and those involved in the communication of ideas between thinkers. He mediates the causal relation between my idea of rain and the bodily processes involved in testifying that 'it's raining', and also the concomitant causal relations between your hearing what I say and the formation of the idea of 'rain' in your mind.[19] We may agree with Hume that this sounds like 'fairy land' (2007, 7.24), but it is important *why* we think occasionalists are wrong. We can agree that their metaphysics and philosophy of mind are mistaken—because, perhaps, there's no reason to believe in God, or to think causation is mediated in such a way, or to accept the idea theory of the mind—but it would seem that Pritchard would also claim that occasionalists are mistaken in their epistemology: knowledge should not be ascribed to us since the occasionalist God gives us a helping hand. This may constitute a novel objection to occasionalism or—and this, I think, is the moral we should draw—it suggests that we should be cautious in accepting Pritchard's claim that hidden helpers always undermine attributions of knowledge.

Fourth, imagine that we discover that there are indeed benevolent angels or omniscient engineers of either divine or alien origin who engineer our beliefs so that they are mostly true. Left to our own devices we would have many more false beliefs than we do. How would such a discovery be described? Pritchard might say something along these lines: 'For millennia we had no idea that our epistemic overlords gave us such a helping hand; we thought we knew a

---

Mark Tebbit for pointing this out. Second, in some ways Iago determines the truth of the belief about the handkerchief. He does not allow Emilia to give the handkerchief back to Desdemona. Othello thus believes the handkerchief is missing, and Iago plots to make it the case that this continues to be so. This therefore approaches a scenario in which, as Lewis puts it, the mind is not passive to the world (see fn. 17). I will ignore this complication.

19 The occasionalist Geulincx talks of how we are mere spectators of our thoughts, those that are caused by God. See J.-C. Bardout (2002, p. 148). Occasionalism is therefore Augustinian in its approach since human cognition and knowledge requires divine assistance or illumination.

great deal about the world, but we were wrong.' I do not think, though, that this is how we would react. This, I claim, is a more plausible response: 'For millennia we had no idea that knowledge worked like *that*—that it involves such a helping hand from our epistemic overlords; it's a good thing they provide such help otherwise we would know very little about the world.' One can also think here of how the children in *Violin* and *Attic* would react when, as they get older, they discover that their beliefs were engineered by their parents. It seems to me that they would discover *how* they knew about the dangers of the attic and their musical prospects, and not that, as it turns out, they did not have knowledge of such things because of their parents' epistemic intervention.

As a child grows older (and wiser) they may come to learn of their parent's intervention and thus come to know *how* they knew about the danger in the attic. They might then be said to have a higher grade of knowledge. They now know, for example, that it's dangerous because of the rickety ladder and because of the heavy, precariously stacked boxes up there (and not because of monsters), and they know that it was their parent's deception that originally led to their having a true belief regarding the dangers of the attic. The younger child lacks such higher-grade knowledge, and this may explain why the child can be seen as epistemically lacking in some way, and consequently it may be thought that such engineering does not result in knowledge. The conclusion to be drawn, though, is only that the child lacks higher-grade knowledge; he still possesses lower-grade knowledge, that is, knowledge that consists of non-accidental true belief that is the result of a cognitive achievement (on the part of the engineer).

Sosa has a related distinction between lower and higher-grade kinds of knowledge—between what he calls animal and reflective knowledge. Animal knowledge is 'apt belief *simpliciter*' (2007, p. 34), true belief that is the result of one's cognitive skills; reflective knowledge is 'apt belief aptly noted' (2007, pp. 34–5), that is, apt belief accompanied by competent assessment that one is in the right conditions to be exercising the relevant competences underlying apt belief. (I say more about Sosa's notion of aptness in §4 below.) Sosa also discusses a case similar to ones that I describe as involving engineering. A jokester controls the colour of a kaleidoscope surface and also the ambient light and thus has two ways of making a particular surface appear red. There's a 'good' combination of red surface and white light, where an observer correctly judges the colour of the surface to be red, and a 'bad' combination of white surface and red light where the observer is misled. Because these cases are indistinguishable from the point of view of the observer, they could be thought to undermine ascriptions of knowledge, even in the good case. Sosa, however, argues that in the good case, '[t]he perceiver would there be said to have apt belief, and animal knowledge, that

the surface is red. What he lacks, we may now add, is *reflective* knowledge, since this requires apt belief that he aptly believes the surface to be red' (2007, p. 32).

I have suggested, then, that the skills of the engineer play a crucial role in the transmission of true beliefs to the hearers in *Iago*, *Violin* and *Attic* and I have attempted to undermine the intuition that knowledge is not attributable in such circumstances even though there are hidden helpers. Such talk of abilities and skills suggests that virtue epistemology might be well-placed to account for such knowledge since on this kind of approach knowledge is seen as a cognitive achievement. The rest of the paper explores this suggestion.

4 Engineered knowledge and virtue epistemology

There are various kinds of virtue epistemology. One such is virtue reliabilism, according to which epistemic virtue amounts to the possession of reliable faculties, those that consistently lead to the acquisition of true belief.[20] Such faculties include those involved in sense perception, inductive and deductive reasoning, and memory. Sosa (2007) offers an analogy with archery to illustrate his version of virtue reliabilism. A lucky gust of wind can blow an arrow onto the bullseye and this is akin to a case where a belief is true by accident. An archer could also consistently hit the bull, but this is because there is a mechanism that moves the target in line with the arrow in flight. This is akin to *Temp*, and to cases of safe belief that are not creditable to the agent. In contrast to these cases, knowledge is akin to an archer hitting the target through his own reliable skill. Knowledge consists in belief that is accurate, adroit and apt. Accuracy concerns whether one hits the target—whether one's beliefs are true. Adroitness concerns whether one's performance manifests a cognitive skill on the part of the believer, and a performance is apt if one's accuracy is due to one's adroitness.

It is difficult, however, to apply this analogy to cases of testimony given that in these there is a role to play for both speaker and hearer, whereas it is only the archer who is responsible for the kind of accuracy definitive of archery. A better analogy is that of playing catch, with the speaker as thrower and the hearer as catcher. A wayward throw could be caught because a gust of wind directs the ball into the catcher's hand. This is lucky. Such a throw could, though, be caught by the catcher pulling off a skilful diving catch. This is analogous to

---

20 See Sosa (2007) and Greco (2010).

a case in which the epistemic work is primarily to the credit of the hearer: a case, perhaps, where a speaker garbles an explanation yet the hearer manages to decipher what they mean. There are also cases where thrower and catcher both play a full part, and Sosa aligns a successful case of testimony to the teamwork behind a touchdown in American football (ibid., p. 94). Further, he considers cases of testimony where credit does not obviously extend to the hearer: not engineered cases, but everyday ones where we take the utterances of others at face value and thereby come to know the time or the direction to the museum merely on the say-so of somebody else. These are the kinds of cases Lackey (2007) takes as undermining virtue epistemology since the hearers of such testimony seem to deserve little or no credit for their true beliefs. Sosa, however, suggests that 'if the correctness of that eventual belief is attributable to a competence, it is not one seated in the believer individually. Any such competence would have to be socially seated instead, in some broader social unit' (2007, p. 93). It is not entirely clear what Sosa has in mind here, but it's plausible to see such competences as the community analogues of reliable cognitive faculties. For example, testimonial chains of beliefs must be initiated by perceptual beliefs, ones acquired via the reliable perceptual capacities of agents several links removed from the knower, and reliable memory must sustain the links in the chain of transmission, aided by accurate written or spoken articulation of the transmitted beliefs and reliable understanding.

In such cases the hearer deserves limited credit—although, perhaps some, given that we are usually selective in deciding who is likely to know the answers to such questions (concerning the time and the museum), and that we wouldn't accept replies that were too outlandish. *Iago*, *Violin* and *Attic*, though, are limit cases in which the credit due to the hearer is close to zero.[21] Engineered knowledge is akin to where almost all of the work is done by the thrower. Think of teaching a child to play catch: everything has to be right about the throw—

---

[21] Even in the limit case the subject exercises her own epistemic abilities to some extent and thus deserves some credit. She must, for example, understand the utterances of the engineer. She would not count as acquiring knowledge if she simply parroted, without understanding, what she was told. Further, her belief must be acquired using capacities that, in normal, non-engineered cases, would contribute to the acquisition of knowledge. She may, for example, monitor the speaker for clues to their reliability, honesty and sincerity; she is not just a sponge soaking up what is said. In the cases described, though, the engineer is careful not to reveal any such clues indicating unreliability.

the spin, the pace, the trajectory; one shouldn't distract the child as she waits for the ball to arrive in her hands; one shouldn't applaud too quickly or the ball may be dropped. Such skills on the part of the thrower are ones that must be learnt, ones that take patience and practice, and ones for which she deserves praise. And, to return to the issue of fragility: on other occasions even such an exemplary throw could be fumbled—the child does not have a *safe* pair of hands—but, nevertheless, on this occasion the ball is caught and this is no accident; it is down to the skills of the thrower. Greco discusses a case in which a player in a game of football scores an easy tap-in goal after being set up by a brilliant pass from another player (Greco, 2010, pp. 82–3). In this case some credit is given to the scorer, even though the passer deserves more credit for the goal. There is a further case, though, and that's where a brilliant pass bounces off the head of a hapless forward and goes in the goal. Here the forward deserves no credit. That may be so, but this still stands as a goal and, if the passer intended the deflection, as an achievement on his part—and this, I claim, is also the case with engineered knowledge.[22]

There are, though, problems in applying the virtue reliabilist approach to engineered knowledge. First, virtue reliabilism does not sufficiently acknowledge the credit deserved by the speaker. In cases of engineered knowledge the engineer deserves praise, and this is true—on epistemic grounds—even in cases that are morally suspect. Second, virtue reliabilism is one-dimensional, only focusing on the virtuous epistemic property of reliability, whereas the skills of deception required of the speaker in cases of engineered knowledge are various. He must, for example, be able to project himself into the place of another—know how they are likely to think given the lies told. Iago must have knowledge of Othello's emotional state and a parent must judge the details of his story so that the child will avoid the attic without running and

---

[22] Pritchard (2010, pp. 36–40) argues that there must be more to knowledge than cognitive achievement since any such achievement can be prey to the 'environmental luck' manifest in fake barn cases. I do not have space to discuss this in detail, but with respect to testimonial exchange we can think of a modified barn case that is analogous to those in which engineering takes place. You may be in Fake Barn County, mostly forming false beliefs about the apparent structures by the road; from time to time, though, your eye settles on a rare real barn and you have a true belief—a belief, though, that is not safe and does not amount to knowledge. However, a friendly barn-spotter (the engineer) sees that you are in an epistemically unfriendly environment and, without explaining what he is doing, points you at one of the real barns. Here, you do have knowledge—that engineered by the barn-spotter.

screaming out the house and without sleepless nights. His timing must be right, as must his body language. He must be creative and imaginative, and know how to adjust his scheme if the audience is resistant. A liar must keep track of his lies; remember what was said and why, since the stories one invents can easily crumble; this is so for villains and parents. The rich sets of skills required in cases of engineered knowledge also therefore favour a virtue responsibilist approach, with various skills and character traits relevant to epistemic success. Epistemic virtues, then, must be acquired characteristics and skills, 'active features of her agency: actions, motivations, and habits over which she has some control and for which she is (to some degree) responsible' (Battaly, 2008, p. 648). The epistemic role of the speaker in engineered cases thus suggests a virtue responsibilist approach.[23]

There is therefore a cognitive achievement on the part of the engineer, and this achievement is due to certain skills that he possesses and that he exercises in these scenarios. There is some tension here in characterising such a performance as involving epistemic *virtue* when deception is involved in this way. This is certainly true in *Iago*, given the morally corrupt purposes to which his deceptions are put, although less so in the cases of *Violin* and *Attic*, cases in which deception is ultimately playing a benevolent role. Such tension can be eased, though, by making sure we do not conflate epistemic and moral considerations. Iago is morally corrupt, but he is epistemically astute and his skills enable him to inculcate true belief (and knowledge) in Othello. It is therefore appropriate to call the skills and character traits involved *epistemic* virtues, although ones exercised in the pursuit of evil. It is, after all, sometimes appropriate to express a kind of admiration for particularly skilful displays of vice. In Homer '[i]t is a measure of wit or *ingenium* to master the art of lying. Not even the gods shun lying and deception, and they model this art for humans' (Weinrich, 2005, p. 64).

Such an approach is sometimes called neo-Aristotelian because of its focus on character traits, but differences to Aristotle's own account should be made clear. For Aristotle, in exercising one's intellectual virtues one is manifesting what is essentially human, thus performing one's proper function and, in doing so, living a fulfilled life. It is, though, problematic to see *Iago* as involving intellectual virtue so construed. Aristotle would see Iago as manifesting epistemic cleverness, a knavish version of practical wisdom. Aristotle also distinguishes virtues from skills. Both may require training and practice and both may ossify if not used: a mnemonic technique to remember the names of one's numerous cousins needs to

---

23 Virtue responsibilist accounts are developed by Code (1987) and Zagzebski (1996).

be practiced, as does one's tactfulness in discussing controversial topics—skill and virtue respectively. One difference between them, it is claimed, is that one acts skilfully solely in virtue of being able to produce something that is excellent in some way. To manifest a virtue, though, one needs to be able to articulate how and why one acted in that particular way. Some of the key traits in my cases would appear to be skills, those such as timing and control of body language; others are virtues, such as a parent's sensitivity in engaging with his children.

Lastly, I shall provide a schema for a kind of virtue responsibilism that would accommodate engineered knowledge. One can think of our epistemic relation to the world as involving layers of intellectual capacities and virtues. First, there are the kind of faculties and competences that play a crucial role in Sosa's virtue reliabilism, and these are indeed necessary in order to have knowledge. One must have good eyesight for perceptual knowledge, and one must have certain linguistic abilities in order to acquire and transmit testimonial knowledge. These faculties, though, must also be used in the right way, and it is virtuous character traits that enable us to do this. One needs to be observant in order that one epistemically benefits from one's good eyesight, focusing one's perceptual capacities when appropriate; intellectual humility and open-mindedness are sometimes required in order that one accepts a testimonial report contrary to one's own beliefs. Such traits are involved in an individual's acquisition of knowledge, and also in the kinds of social exchanges characteristic of engineered knowledge, with, as said, the relevant virtues here being those of the engineer. There is also a third level of epistemic engagement: at times an analogue of Aristotle's practical wisdom is required in order to guide the exercise of a variety of virtues on particular occasions. As a parent one knows that honesty is to be inculcated in one's children, but one also knows the power and usefulness of deception. These are competing considerations that must be judged in particular circumstances. It may be acceptable to say that there are monsters in the attic in order to protect the child from danger, but not to say that the monsters will come if homework is not completed.

Traditionally, lines of knowledge transmission are clean—that is, they are seen as involving testimony that is sincere and reliable, and they are driven by emotions and traits such as wonder at the detail of nature, honesty and curiosity.[24] I have suggested, though, that the art

---

[24] Morton (2014) subverts discussion of epistemic emotions and traits by noting that epistemic vices can also play a role in knowledge acquisition, even claiming that 'the search for hypotheses and evidence may go better if individual epistemic agents are moved by less worthy emotions' (ibid., p. 165). He discusses the role played by group loyalty, antipathy [and] rivalry'

of lying can also play an epistemic role for parents and Shakespearean plotters. These are, of course, unusual kinds of cases, but perhaps they shed light on testimonial transmission in general by allowing us to see through certain prevalent claims concerning necessary features of knowledge transmission: it need not be sincere, it can be fragile, and it can depend on the epistemic role of the speaker—features that may also be relevant to everyday, non-devious cases of knowledge transmission.[25]

---

and is 'interested in the role emotions that we do not easily think of as admirable can play in motivating epistemic virtues' with the result that 'some virtuous combinations of vices result in more knowledge . . . than we can have from the enterprises of dispassionate sages' (ibid., p. 171).

[25] Acknowledgements.

References

Audi, R. (2013). Testimony as a social foundation of knowledge. *Philosophy and Phenomenological Research, 87*(3), 507–31.

Bach, K. (1994). Conversational implicature. *Mind and Language 9*, 124–62.

Bardout, J.-C. (2002). Occasionalism: La Forge, Cordemoy, Geulincx. In S. Nadler (Ed.), *A companion to early modern philosophy* (pp. 140–51). Oxford: Blackwell.

Battaly, H. (2008). Virtue epistemology. *Philosophy Compass*, *3*(4), 639–63.

Bobier, C. (2014). Anti-luck virtue epistemology and divine revelation. *Philosophia*, *42*, 309–20.

Bok, S. (1999). *Lying: moral choice in public and private life*. New York: Vintage Books.

Carston, R. (2002). *Thoughts and utterances*. Oxford: Blackwell.

Chabon, M. (2013). The crying of September 11. Review of T. Pynchon, *Bleeding edge. New York Review of Books.* http://www.nybooks.com/articles/2013/11/07/thomas-pynchoncrying-september-11/. (Accessed 20 April 2017)

Code, L. (1987). *Epistemic responsibility*. Providence: Brown University Press.

Comesana, J. (2005). Unsafe knowledge. *Synthese*, *146*, 395–404.

Dworkin, G. (2014). The care and management of lies. http://www.3quarksdaily.com/3quarksdaily/2014/12/the-care-and-management-oflies.html. 8.12.2014. (Accessed 20 April 2017)

Fitelson, B. (2010). Strengthening the case for knowledge from falsehood. *Analysis*, *70*, 666-69.

Fricker, E. (2006). Second-hand knowledge. *Philosophy and Phenomenological Research*, *LXXIII*(3), 592–618.

Gaultier, B. (2014). Achievement, safety and environmental epistemic luck. *Dialectica*, *68*(4), 477–97.

Goldberg, S. (2005). Testimonial knowledge through unsafe testimony. *Analysis*, *65*(288), 302–11.

Graham, P. (2014). Against transglobal reliabilism. *Philosophical Studies*, *169*(3), 525–35.

Greco, J. (2003). Knowledge as credit for true belief. In M. DePaul & L. Zagzebski (Eds.), *Intellectual virtue: perspectives from ethics and epistemology* (pp. 111–34). Oxford: Oxford University Press.

Greco, J. (2010). *Achieving knowledge: a virtue theoretic account of epistemic normativity*. Cambridge: Cambridge University Press.

Grice, P. (1989). *Studies in the ways of words*. Cambridge, MA: Harvard University Press.

Hetherington, S. (1999). Knowing failably. *Journal of Philosophy*, *96*(11), 565–87.

Hume, D. (2007). *An enquiry concerning human understanding*. T.L. Beauchamp (Ed.). Oxford: Oxford University Press.

Klein, P. (2008). Useful false beliefs. In Q. Smith (Ed.), *Epistemology: new essays* (pp. 25–62). Oxford: Oxford University Press.

Lackey, J. (2008). *Learning from words: testimony as a source of knowledge*. Oxford: Oxford University Press.

Lackey, J. (2007). Why we don't deserve credit for everything we know. *Synthese*, *158*, 345–61.

Levy, N. (2014). Frankfurt in fake barn country. *Metaphilosophy*, *45*(4–5), 529–42.

Lewis, D. (1973). The thermometer model of knowledge. In S. Bernecker and F. Dretske (Eds.), *Knowledge*: *readings in contemporary epistemology* (pp. 72–85). Oxford: Oxford University Press.

Malebranche, N. (1997). *The search after truth*. T. Lennon & P. Olscamp (Eds.). Cambridge: Cambridge University Press.

McCain, K. (2014). Testimonial knowledge from lies. *Philosophia*, *42*(2), 459–68.

Morton, A. (2014). Shared knowledge from individual vice: the role of unworthy epistemic emotions. *Philosophical Inquiries 2*(1), 163–72.

Nozick, R. (1981). *Philosophical explanations*. Cambridge, MA: Harvard University Press.

Nyberg, D. (1993). *The varnished truth: truth telling and deceiving in ordinary life*. Chicago: University of Chicago Press.

Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.

Pritchard, D. (2006). A defence of quasi-reductionism in the epistemology of testimony. *Philosophica*, *78*, 13–28.

Pritchard, D. (2009). Safety-based epistemology: whither now?. *Journal of Philosophical Research*, *34*, 33–45.

Pritchard, D. (2010). Anti-luck virtue epistemology. In D. Pritchard, A. Millar & A. Haddock, *The nature and value of knowledge: three investigations* (pp. 48–65). Oxford: Oxford University Press.

Pritchard, D. (2012a). Anti-luck virtue epistemology. *Journal of Philosophy*, *109,* 247–279.

Pritchard, D. (2012b). Apt performance and epistemic value. In J. Greco & J. Turri (Eds.) *Virtue epistemology: contemporary readings* (pp. 117–29). Cambridge, MA: MIT Press. (Originally published in *Philosophical Studies*, *143* (2009), 407–16.)

Riggs, W. (2012). Why epistemologists are so down on their luck. In J. Greco & J. Turri (Eds.) *Virtue epistemology: contemporary readings* (pp. 285–305). Cambridge, MA: MIT Press. (Originally published in *Synthese*, *158* (2007), 329–44.)

Sosa, E. (1991). *Knowledge in Perspective*. Cambridge: Cambridge University Press.

Sosa, E. (1999). How to defeat opposition to Moore. In J.E. Tomberlin (Ed.) *Philosophical Perspectives*, *13*, *Epistemology*, 141–53.

Sosa E. (2007). *A virtue epistemology*. Oxford: Oxford University Press.

Turri, J. (2012). Is knowledge justified true belief? *Synthese*, 184, 247–59.

Warfield, T. (2005). Knowledge from falsehood. *Philosophical Perspectives*, *19*, 405–16.

Weinrich, H. (2005). *The linguistics of lying and other essays*. Seattle: University of Washington Press.

Welbourne, M. (2001). *Knowledge.* Chesham: Acumen.

Williamson, T. (2000). *Knowledge and its limits.* Oxford: Oxford University Press.

Zagzebski, L. (1994). The Inescapability of Gettier Problems. *The Philosophical Quarterly*, 44(174), 65–73.

Zagzebski, L. (1996). *Virtues of the mind: an inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge: Cambridge University Press.