

Understanding, Luck, and Communicative Value

Andrew Peet, University of Leeds.

*Penultimate draft for A. Fairweather & C. Montemayor (eds.) **Linguistic Luck** (Forthcoming, OUP). Please cite published version.*

1 Utterance Understanding

‘Understanding’ has several different uses. Firstly, we can talk of understanding a concept or a topic. This plausibly requires knowing the answer to a certain why question, or perhaps possessing a web of interconnected beliefs regarding a topic. This is not the sense of ‘understanding’ we will be focusing on here. Secondly, we can talk of understanding a sentence. For example, I might understand the sentence ‘Jeg er ikke her’ (it means ‘I am not here’). This is closer to our topic. Thirdly, we can talk of understanding what is said. That is, of understanding somebody’s utterance. It is this third notion of understanding upon which we shall focus.

One can plausibly understand the sentence somebody has used without understanding their utterance. For example, suppose I get home to find a note with the words ‘Jeg er ikke her’ written on it. I understand the sentence. However, since I do not know who wrote the note, I do not understand the utterance. I cannot grasp what it is saying. Suppose it was written by my wife, but I take it to have been written by my son. Then I have misunderstood the utterance despite understanding the sentence. Consider another example: I am out for dinner with a friend and her steak is undercooked. She complains that her steak is ‘raw’. Suppose I take her literally and think she means that her steak is completely uncooked. Then although I have understood the sentence she has used, I have not understood her utterance.

Not only is sentence understanding insufficient for utterance understanding, it is also unnecessary. Suppose I have just had a carpentry accident and managed to get a large splinter stuck in my eye. I am considering whether or not to seek treatment, so I speak to my sister who is a medical professional. She tells me that ‘No eye injury is too small to be ignored’. She intends to communicate that no eye injury is small enough that it can be ignored: you should always get an eye injury examined. We both take this to be the literal meaning of the sentence she has uttered. However, we are both wrong. The sentence ‘No eye injury is too small to be ignored’ really says that no matter how small the eye injury is it can be ignored. So, neither myself nor my sister understand the sentence she has uttered. Yet I plausibly understand her utterance.

Strictly speaking the proposition the speaker intends to directly communicate with an utterance can differ from the proposition they actually express. So, we could further distinguish two forms of utterance understanding - one corresponding to recovery of the speaker’s intended proposition, and another corresponding to recovery of the truth conditional content of their utterance. However, in most cases truth-conditional content will correspond to the proposition the speaker aims to directly communicate with their utterance. So, typically these forms of

understanding will coincide.¹ My focus here will be on recovery of the speaker's intended proposition. So, I will use 'what is said' to denote the proposition the speaker intends to directly communicate via their utterance. It is also worth noting that understanding an utterance will involve more than recovering the proposition intended by the speaker. It will also involve correctly judging the force of the utterance (i.e. is it an assertion? a suggestion? a sarcastic remark?). I will ignore this complication in what follows. Similar issues to those that arise for content will arise for force, and similar solutions will be applicable.

2 Understanding and Reliability

I will be exploring the prospects for reliability-based approaches to utterance understanding. That is, I will be considering approaches committed to some version of the following:

Reliable Recovery: Understanding requires reliable recovery of what is said.

Many theorists go beyond this, holding that utterance understanding requires knowledge of what is said.² However, as we will see, it is hard enough to maintain a mere reliability condition on understanding, let alone a knowledge condition. So, the prospects for knowledge views look dim.³

Before going through the arguments for **Reliable Recovery**, it will be helpful to distinguish between two different senses of 'reliable': local and non- local reliability. The distinction between local and non-local reliability comes from epistemology, and corresponds to different ways in which a belief might be reliable. A true belief is locally reliable if, given the way that very belief was formed, it was not a matter of luck that it was true. Local reliability conditions are often captured in modal terms. For example, it may be thought that for a belief to be locally reliable it must be safe:

Safety: A belief is safe iff there are no nearby worlds in which it is formed in the same way and false.⁴

A belief is non-locally reliable if it is formed in a manner that would yield true beliefs in some range of specific circumstances (that may or may not include the actual circumstances of formation). For example, we might say a belief is non-locally reliable if it is formed in a way that

¹ See Peet (forthcoming) for an account of the relationship between intended and truth conditional content

² Examples include Campbell (1982), Dummett (1978, 1991), Davies (1989), Evans (1982), Heck (1995), Higginbotham (1992), and McDowell (1994).

³ For arguments against the knowledge view see Fricker (2003), Hunter (1998), Longworth (2008, 2009, 2018), and Pettit (2002). And for arguments against the related claim that testimonial knowledge requires knowledge of what is said see Peet (2018a, 2019). Many of the cases discussed in this paper are problematic for the knowledge view.

⁴ **Safety** is widely endorsed, but for influential defenses see Sosa (1999), Williamson (2000), and Pritchard (2005). I will assume that **Safety** correctly captures local reliability in what follows, although little turns on this.

yields true beliefs in normal circumstances.⁵ It should be clear that a belief can be locally reliable without being non-locally reliable, and vice versa. Local reliability is typically thought to be necessary for knowledge. Non-local reliability conditions typically feature in theories of justification.

With this said, it should be clear that **Reliable Recovery** is ambiguous. Separate principles could be formulated with respect to local and non-local reliability (or, both could be required). We can get clearer on the precise sense in which utterance understanding requires reliability by considering the arguments for **Reliable Recovery**.

There are a number of potential arguments for **Reliable Recovery**. The first two arguments draw on Longworth's (2018) discussion of the knowledge of what is said condition. As Longworth demonstrates, neither argument provides convincing support for a knowledge condition on understanding. However, they have greater force in support of **Reliable Recovery**.

The first argument starts from the assumption that utterance understanding regularly yields knowledge of what is said. We would typically infer from the fact that somebody understands another's utterance to the conclusion that they know what the speaker has said. If we are to explain the naturalness of this inference we must endorse something like the following: usually, or in normal circumstances, understanding yields knowledge of what is said. This tells us two things. Firstly, it tells us that in normal circumstances understanding yields a veridical judgement about what is said. That is, normally when we understand somebody to have said that p they will have said that p. So, this already gives us a form of non-local reliability. However, we can go further. When we gain knowledge that somebody has said that p our judgment must also be locally reliable. That is, it must satisfy an anti-luck condition such as the safety condition. So, we also learn that in normal conditions utterance understanding yields a locally reliable judgment about what is said.

The second argument is similar but starts from the observation that utterance understanding serves as a basis for the acquisition of testimonial knowledge. Plausibly, recovery of what is said is essential for testimonial knowledge. If I take you to have said that p, and thereby come to believe p, then unless you have actually said that p I cannot gain knowledge that p from your testimony.⁶ So, if understanding underwrites the acquisition of testimonial knowledge (i.e. leads to testimonial knowledge in most cases), it must usually yield an accurate judgement about what is said. Moreover, if my recovery of what is said is merely lucky, then this will not be sufficient for knowledge, as my testimonial belief will be Gettiered. Therefore, understanding must yield a safe judgement regarding what is said in the majority of cases.

⁵ See Goldman (1986), Leplin (2007), Smith (2016), and Graham (2016, 2017) for versions of this approach.

⁶ Strictly speaking I believe this is false. Testimonial knowledge does not require recovery of what is said or intended. Rather it requires non-lucky recovery of a sufficiently similar proposition. Or, more precisely, it requires that the audience non-accidentally recover a proposition that matches the speaker's intended proposition in truth value. I spell out this view in detail in my (2019). For an alternative approach to the communicative preconditions for the acquisition of testimonial knowledge that also denies that the speaker and hearer must entertain the same proposition see Pollock (Forthcoming). For the sake of simplicity, I will stick with the simpler 'recovery of what is said' condition here. The main points generalize to the less idealized approaches just mentioned

The two arguments just presented suggest that understanding is reliable in the following ways: Firstly, understanding usually, or in normal circumstances, yields a veridical judgement regarding what is said. Secondly, it usually yields a safe (i.e. non-lucky) judgement about what is said. However, I believe we can go further. There is reason to believe that understanding always yields a non-lucky judgement regarding what is said. There are cases in which an audience fails to understand an utterance despite reaching an accurate judgement about what is said. This failure of understanding is seemingly due to the luckiness of the audience's recovery of what is said. The classic example is presented by Loar (1979), so we will call such cases 'Loar cases':

Loar: Suppose that Smith and Jones are unaware that the man being interviewed on the television is someone they see on the train every morning, and about whom in that latter role, they have just been talking. Smith utters 'He is a stockbroker' intending to refer to the man on the television; Jones takes Smith to be referring to the man on the train. Now Jones, as it happens, has correctly identified Smith's referent, since the man on the television is the man on the train; but he has failed to understand Smith's utterance. Loar 1976: 357.

Assuming that 'he' is directly referential, the proposition entertained by Smith is identical to the proposition entertained by Jones. Yet Jones has intuitively failed to understand Smith's utterance. Loar concludes that 'he' must not be directly referential. Instead, he holds that interlocutors must coordinate not only on referents but also on modes of presentation in order for communication to yield understanding. However, there is reason to be suspicious of this diagnosis. As Byrne and Thau (1996) note, we can generate similar cases in which the interlocutors coordinate on descriptive modes of presentation. They give the following example:

Hospital A patient checks into hospital and is assigned room 101. Tony dubs him "Winston" and the cognitive value she attaches to the name is: the amnesiac in room 101. Alex is thoroughly unaware that Tony has seen the patient, but by sheer chance she also dubs him "Winston" and attaches the same cognitive value to the name. Alex utters "Winston will never recover" in Tony's presence, and Tony forms the belief she would express by saying "Winston will never recover". Byrne and Thau (1996), p 147.

Indeed, as I note in my (2019), we can generate similar cases by having interlocutors luckily coordinate on, for example, quantifier domains, the meanings of general terms, or the modal bases for modal claims. In general, understanding will be lacking whenever recovery of what is said occurs in a lucky way. This suggests that understanding always yields a non-lucky judgement regarding what is said. When the audience stumbles upon what is said by luck, they will not have understood the utterance.

So, the considerations adduced so far suggest that there is a local reliability condition on understanding. This is how **Reliable Recovery** will be understood for the time being. In the next section I will consider some counter examples to the necessity of local reliability for utterance understanding. These challenges force us to take a more nuanced approach when formulating our local reliability condition.

3 Counter Examples to Local Reliability

Despite being well motivated, reliable recovery seems open to counter example. That is, there seem to be cases of understanding in which reliable recovery fails. This suggests that local reliability is not necessary for understanding. I will consider two forms of counter example here. The first derives from my (2018). It was originally intended as an example of testimonial knowledge without knowledge of what is said. But it also works as a case of understanding without knowledge that ‘S said that p’. The case is as follows:

Mad Scientist The philosophical mad scientist is at it again. His victim is Sally, a car enthusiast. This time, instead of envying his victim, he has implanted a special chip in her brain. This chip causes her to sometimes say ‘that is a fuel-efficient car’, but only when she is in the presence of fuel-efficient cars, and only when she does not intend to do so. It works as follows: whenever Sally is in the presence of a fuel-efficient car it turns on and randomly selects one of two values. If it selects value 1 it switches off and becomes inactive again. However, if it selects value 2 it has the following effect: If Sally doesn’t intend, and doesn’t gain the intention, to comment on the fuel efficiency of the car, then it forces her to utter the sentence ‘that is a fuel-efficient car’. One day Sally and Matt are walking through the city when Sally sees a particularly fuel-efficient car. She considers commenting on its fuel efficiency but hesitates because she doesn’t know if Matt has any interest in cars. She decides on a whim to just go for it and says ‘that is a fuel efficient car!’. On the basis of Sally’s assertion Matt forms the true belief that the car is fuel-efficient. What neither Sally nor Matt know is that this was a case in which the chip selected value 2, so Sally would have uttered ‘that car is fuel-efficient’ even if she did not intend to. Peet 2018, 69.

Two things seem clear: Firstly, Matt understands Sally’s utterance. This is highly intuitive. Moreover, he is employing his normal abilities of comprehension, he gained these abilities in a reliable way, the output is fully veridical, and the causal chain leading to his judgment is non-deviant. Secondly though, there is an important sense in which his recovery of what was said is unreliable: his judgment that Sally said that the car is fuel efficient (henceforth ‘p’) was unsafe.

There are two obvious responses available to the defender of reliable recovery. Firstly, it might be suggested that his ‘Sally said that p’ judgment was actually safe. That is, it might be claimed that in all nearby worlds Sally did say that the car is fuel efficient. However, in many of these worlds she only said it because the chip in her brain made her say it. Alternatively, it may be suggested that understanding only requires a reliable judgment that ‘it has been said that p’. In this case it might be maintained that in nearby worlds Sally does not say that p, but the chip does say that p using Sally as a mouthpiece.

I don’t find either response compelling. Saying is necessarily an intentional action. Whilst it is clearly possible that we can say a particular thing by accident (i.e. we can accidentally say p when we mean to say q), it is doubtful that our having said anything at all could be accidental. For example, if I am asleep or under hypnosis, and words escape my mouth without any intention or awareness on my behalf, it seems wrong to claim that I have said something, even if the sounds I

produce resemble a meaningful sentence. When the chip sends electrical signals through Sally's brain this cause her mouth to open and certain worlds to spill out. She has no intention or motivation to make the relevant noises, and the production of the noises is not under her control or guidance. The saying is, thus, in no way attributable to her. Moreover, it makes little sense to claim that the chip says p. The chip is an inanimate object. It can't do things. This can be made even clearer if we alter the case a little and maintain that that the chip gained its properties randomly, not through design.

A more promising response would be to modify our modal anti-luck condition so that we only consider nearby worlds at which the speaker said something. That is, we could require that there be no nearby worlds in which both A) the speaker produced an utterance, and B) the hearer's judgment that 'S said that p' was mistaken. This allows us to deal with Mad Scientist: After all, there are no nearby worlds in which Sally says 'that car is fuel efficient' and Max fails to recover this proposition. In the few nearby worlds in which she does say that the car is fuel efficient, he reaches the correct judgement.

Unfortunately, this is still not satisfying. Firstly, As Longworth (2008, 2018) notes, young children are able to understand speech. Yet they may not possess the concept of 'saying that'. The requirement that understanding involve a judgment that 'S said that p' seems to over-intellectualize understanding. So, as formulated, our anti-luck condition seems overly demanding.⁷ Instead, we can follow Longworth and hold that understanding what is said merely involves the audience's entertaining of the proposition the speaker intends to communicate. In this case we might spell out reliable recovery in terms of it being non-lucky that the proposition entertained by the audience corresponded to the proposition intended by the speaker. That is, we could require merely that there be no nearby worlds in which both A) the speaker produced an utterance, and B) the proposition entertained by the audience delivered from the proposition intended by the speaker (for ease of exposition I will continue to speak of 'judgements about what is said'.⁸

Unfortunately, this view is also open to counter example.⁹ Consider the following case from Pettit (2002):

German: Imagine that you are travelling in Germany. You are a moderately competent speaker of German, but you come across an unfamiliar word, say, the word 'Krankenschwester'. You see a kindly-looking, elderly German sitting on a bench nearby, and you ask him what 'Krankenschwester' means, hoping that he might know some English. With an air of authority, he smiles and politely replies in English 'it means nurse

⁷ It may be possible to avoid this problem by only requiring a reliable ability to carve up the space of possibilities in the same way that 'S said that p' carves up the space of possibilities. I will put this possibility to one side in what follows.

⁸ This is similar to the condition presented in my (2019) where I require (for the related condition of knowledge-yielding communication) that there be no nearby world in which both A) the speaker said something, and B) the propositions intended by the speaker and recovered by the hearer differ in truth value (with some additional caveats).

⁹ It is actually open to counterexample along multiple lines. As I point out in my (2019), spelling out the anti-luck condition in purely modal terms also leads to problems as we can generate cases of modally stable luck (also see Lackey (2008), and Broncano-Berrocal (2018)). Hyska (MS) raises a similar worry for modal accounts of communicative luck. I resolve this issue by building a 'no-coincidence' clause into the account. I will set this complication to one side here

'which is indeed what the word 'Kranken- schwester' means. Satisfied with his answer, you thank him and go on your way. As a result of this exchange, you are now able to use this previously unfamiliar word correctly and correctly interpret it as it is used by other speakers of German. If a German speaker assertively utters the sentence 'Die Krankenschwester ist nett' for example, you will correctly take the speaker to be asserting that the nurse is nice. Or if you want to say in German that the nurse is coming, you will correctly express this thought with the sentence 'Krankenschwester kommt'. In short, in a familiar sort of way, you have come to understand the word 'Krankenschwester'. However, suppose that, unbeknownst to you, the elderly gentleman-call him Herr Verrückt - is quite senile and doesn't know a word of English. His reply to your question (namely, 'it means nurse') is something he once overheard, but he has no idea what it means or what he is saying when he utters it. In his senility, he has taken to repeating this to tourists, regardless of what he is asked. By sheer coincidence, this was the right answer to the question you happened to ask him. But, had you asked him the location of the nearest post office, he would have said the very same thing. Pettit, 2002, 519-520.

Suppose I am told that 'Die Krankenschwester ist nett'. I thereby have an experience as of the speaker asserting the proposition that the nurse is nice. Plausibly, I understand the utterance. And there is certainly a sense in which my understanding is reliable: in normal circumstances the experience of hearing 'Die Krankenschwester ist nett' and thereby coming to entertain the proposition that the nurse is nice will yield a veridical judgement about what is said. However, it is far less obvious that my judgement is locally reliable. That is, there is a clear element of luck in my recovering the correct proposition. It is, we might imagine, purely by luck that the old man repeats 'it means nurse' to tourists. He could easily have picked some other sentence to repeat such as 'it means chair'. If this had happened, then I would have experienced the speaker as asserting the proposition that the chair is nice. Since this easily could have happened, I don't safely recover what is said. I seem to understand the utterance, yet my recovery of what is said is lucky.¹⁰

In light of this, we may simply retreat to the claim that understanding requires non-local reliability. However, we are then left with the problem of explaining why understanding is absent in Loar cases. Thus, we have a puzzle. Loar cases seemingly demand a local reliability condition for understanding. Yet cases like **German** seemingly show the opposite.

4 A Solution

We must identify a form of luck according to which the audience's judgements in **Loar** and **Hospital** are lucky, but the audience's judgement in **German** is not. To do this we must note a familiar feature of epistemic anti-luck conditions: they are almost always relativized to a method of belief formation. For example, as noted earlier, safety conditions on knowledge hold that an agent's belief, if it is to constitute knowledge, must be true at all nearby worlds in which it is formed in the same way. Without this relativization to methods we would struggle with cases like the following:

¹⁰ Pettit explicitly focuses on understanding a word or a sentence, rather than understanding an utterance. His claim is that, in this situation, one understands the sentence 'Die Krankenschwester ist nett'. However, it is also plausible that one understands literal utterances of this sentence.

Old Woman Lynn is an old woman. She has a son with severe medical issues: he is often close to dying. He comes to visit her and, upon seeing him in front of her, she forms the belief ‘my son is alive’. However, if her son was dead then her younger sister would have come and told her that he was alive. So, she would have believed he was alive even if he was dead. C.f. Nozick (1983).

There are nearby worlds in which the old woman’s son is dead. And in these worlds, she would still believe that he is alive. However, her belief is clearly not lucky. After all, she can see her son standing right in front of her. Relativizing the safety principle to methods of belief formation resolves this issue: there are no nearby worlds in which her belief is false and formed in the same way.

The key to resolving our puzzle, I suggest, rests in the relativization of our anti-luck principle. The basic idea is that we must hold fixed the ‘interpretative base’ of the audience’s judgement regarding what is said (Peet (2019)). That is, understanding will fail whenever there are nearby worlds in which some aspect of the linguistic or interpretative basis of the audience’s judgement remains fixed and results in a failure of coordination with the speaker.

It is clear how this deals with Loar cases: in **Loar**, had the man on the train not been the man on the TV, and had Jones reached his judgement in the same way (i.e. relied on the same meaning schemas, drawn the same contextual inferences etc.), the proposition Jones recovered would have differed from the proposition Smith intended.

However, it is less obvious how this deals with **German**. After all, does the protagonist’s trust in the old German man not constitute part of the interpretative base for their judgement? Is the assumption that the old German man was speaking truthfully not part of the basis upon which they reach their judgement? If so, then our anti-luck condition will fail in **German** as well. And that would be problematic. After all, **German** appears to be a case of understanding.

The solution is to carefully delimit exactly which factors are to be included in the interpretative base to which we relativize our anti-luck condition. Should we only include the immediate information drawn upon by the hearer – i.e. information about word meaning, and judgements about context etc? Or should we go further back, and include information about how these judgements regarding context and word meaning were reached?

When we are thinking about knowledge-yielding communication (i.e. the communicative preconditions for the acquisition of testimonial knowledge) it is quite clear that we ought to factor in this wider information. The protagonist in **German** does not, and cannot gain testimonial knowledge when they are told, for example, that ‘Die Krankenschwester ist nett’. Their belief will be luckily true. This is because the basis upon which they acquired their beliefs about word meaning (or, their schemas for application of word meaning in context), was unreliable.

However, as **German** shows, there can be cases of understanding without knowledge-yielding communication. So, it is not obvious that we must factor in this information when assessing whether a hearer has understood an utterance. Rather, I suggest that understanding has a weaker local anti-luck condition: there must be no nearby worlds in which the speaker produces a

relevantly similar utterance and audience's judgement has the same 'direct basis' in which they fail to recover what is said. The direct basis of an audience's judgement regarding what is said will include factors such as their beliefs (or schemas) regarding word meaning, together with the beliefs about context upon which their judgement (i.e. their state of comprehension) is based. It will not include the processes that gave rise to these beliefs about context and word meaning.

In **Loar** and **Hospital** even this weak local anti-luck condition fails: even if we just hold fixed the direct basis of the hearer's judgements, they could easily have been led astray. However, this weak local anti-luck condition is satisfied in **German**: the direct basis for the hearer's judgement did not include any information about the old man. It merely included beliefs such as 'Krankenschwester' means nurse. There were no nearby worlds at which the hearer reached their judgement on such a basis and nonetheless failed to coordinate with the speaker. Thus, I suggest we can resolve our puzzle by requiring that understanding merely satisfy this weaker anti-luck condition.

It is important to note that this form of local reliability is too weak, by itself, to support knowledge of what is said, or testimonial knowledge. When our protagonist first heard 'Die Krankenschwester ist nett' they understood what was said. However, they were not in a position to know what was said and were not in a position to gain testimonial knowledge. So, the considerations adduced so far suggest the following:

1. In normal circumstances understanding yields a safe veridical judgement about what is said – a judgement reliable enough to support knowledge of what is said, and testimonial knowledge.
2. In all circumstances understanding satisfies a weaker local anti-luck condition. However, this condition is not by itself strong enough to support knowledge of what is said, or testimonial knowledge.

This view fits the data nicely. However, it may appear somewhat ad hoc. It is natural to wonder at this point why understanding should have an anti-luck condition, especially the specific anti-luck condition I have outlined here. This worry naturally leads on to a deeper challenge for **Reliable Recovery**.

5 Luck and the Value of Understanding

There is a long standing question in epistemology regarding the value of knowledge. Traditionally this problem has been thought of in terms of the value of knowledge over true belief. However, since the emergence of the Gettier problem, focus has shifted to the question of what makes knowledge more valuable than justified true belief. That is, what does the reliability or anti-luck condition add to knowledge? Why is it better to be in a state that satisfies the anti-luck condition than an otherwise identical state that does not?

Linda Zagzebski (2003) nicely illustrates the problem with the following analogy: Suppose we have two coffee machines. One reliably produces great coffee. The other hardly ever produces great coffee. However, on this one occasion, both machines have produced identical cups of coffee.

Is the cup of coffee produced by the reliable coffee machine better, just because of the reliability of its source, than the otherwise great cup of coffee luckily produced by the unreliable coffee machine? It is not clear why it would be. So, why should we consider a reliably produced justified true belief to be any more valuable than an unreliably produced justified true belief?

Megan Hyska (2018, MS) has recently argued that a similar problem arises for understanding. Hyska agrees that understanding must satisfy a local anti-luck condition. However, she observes, it is not clear what value this anti-luck condition could add over and above the other conditions on understanding. Other than the reliability requirement, Hyska (2018) thinks of understanding in a broadly Gricean manner, formulated as follows:

An audience understands an utterance (or intentional signal) that p just in case:

1. she entertains that p
2. (a) she entertains that p because of the signal, and
(b) it is manifest to her that the speaker intended the signal to cause her to entertain that p.
3. p is (consistent with) the content the speaker intended to get across with the signal.

Hyska (2018) 56-57. C.f. Grice (1957).

Why, she asks, would it be better to satisfy these conditions in a reliable or non-lucky way, than it would be to do so in an unreliable or lucky way? What does reliability add to the end product?

It might seem that there is a clear answer to Hyska's worry: as observed from the outset, understanding typically supports the acquisition of testimonial knowledge. Usually, when one understands an utterance, this will put one in a position to gain testimonial knowledge (as long as the speaker is reliable and sincere etc.). If understanding was not necessarily reliable, then it would not be capable of supporting the acquisition of testimonial knowledge. So, assuming that knowledge is valuable, we should be able to resolve the value challenge for understanding.

However, the matter is not so simple. Firstly, as we have already seen, although understanding must be locally reliable, satisfaction of this local reliability condition is not sufficient for knowledge (although it is necessary). Secondly, this answer would render the value of understanding purely epistemic. Yet, Hyska would argue, understanding is a 'signaling achievement' – a communicative achievement, not a purely epistemic achievement.¹¹

It is clear that communication often yields knowledge. This is part of why communication matters to us. The spread of knowledge is, in an important sense, one of the functions of communication. However, it is not clear that the *fundamental* goal – the telos of communication – is knowledge acquisition. After all, communication can clearly succeed without knowledge being

¹¹ Hyska (2018, p 78) suggests that the signaling domain is the 'proprietary domain of the value of communication'. She argues that it is, like the epistemic domain, goal directed. But unlike the epistemic domain, the goal is not truth, but rather information transfer.

gained. This is what we saw with **German**: communication was successful, yet the audience was neither in a position to gain testimonial knowledge, nor to gain knowledge of what is said.

This is not the place to settle the question of the fundamental goal of communication.¹² I don't think we have to answer this question in order to answer the question of what value is added by reliability.

6 Luck, Credit, and Value

One of the most promising responses to the epistemic value problem comes from virtue epistemology.¹³ Virtue epistemologists suggest that knowledge is an achievement – it is something for which we deserve credit. The basic idea is that we can normatively appraise goal directed performances in the following ways:

Success: Does the activity achieve its goal?

Competence: Is the performance produced by an ability to achieve success relative to the goal?

Aptness: Is the performance successful in virtue of the agent's competence? C.f. Sosa 2007.

Ability is standardly thought of as requiring a reliable disposition to succeed at the task in hand. It can also be thought of in terms of the possession of a way of achieving the end that, if employed, would reliably yield success. An agent deserves credit for their success only when their performance is apt. Consider archery: the goal of archery is to hit the target. However, one can achieve success at this practice without deserving credit for it:

Beginner Luck: Suppose a first time archer with absolutely no skill or knowledge of archery hits the target. They succeed in their aim. However, their success is due to luck.

Skilled Luck: Suppose a highly skilled archer fires at the target. Their performance is competent: it is produced by a reliable ability to hit the target in normal conditions. However, conditions are not normal. First a gust of wind blows the arrow off target. Then another surprise gust of wind blows it back on target. As a result, the target is struck. Although the archer succeeded in hitting the target, and although their performance was competent, their success was due to luck, as their success was not explained by their ability. (C.f. Sosa (2007))

In both cases the goal was achieved in a lucky way. As a result, the performances were not apt, and the agents don't deserve credit for their success. Had the second archer's success been

¹² Although I am sympathetic to Jessica Keiser's (Forthcoming) suggestion that the goal of communication is to direct attention. Hyska's suggestion is that communication aims at the transfer of information.

¹³ See Sosa (1991, 2011, 2015), Greco (1999, 2000, 2010), Zagzebski (1996), Kelp (2011, 2016, 2019).

explained by their ability at archery then their success would not have been lucky, and their shot would have been creditworthy.

Skilled Luck is analogous in structure to standard Gettier cases. In Gettier cases an agent competently forms a true belief. The practice of believing aims at truth. So, the agent is successful. However, their success is not due to their competence, so they do not deserve credit for their success.

This provides a natural response to the value problem for knowledge: knowledge is better than Gettiered true belief because, when a belief is Gettiered, the believer does not deserve credit for their success. Perhaps we could think of the value of understanding similarly? Here are the ways we might assess an audience in a communicative exchange:

Success: Has the audience successfully recovered the intended message?

Competence: Is their state of entertaining a particular proposition produced by an ability to recover intended messages.

Aptness: Is the audience's recovery of the intended message explained by their ability to recover intended messages?

By maintaining that understanding must be apt we are able to explain the value of understanding: it is an achievement, something for which we deserve credit. We are also able to explain why understanding is incompatible with luck: creditworthiness is inconsistent with luck. Moreover, we are able to explain these matters without making the value of understanding epistemic. Understanding can be thought of as an achievement in the signaling domain.

Unfortunately, the credit approach to the value of understanding faces some major challenges. In order to deal with standard Gettier cases in Epistemology the credit theorist has to maintain that the subject's cognitive abilities are the most salient explanatory factor in producing their success (Greco (2003)). After all, in typical Gettier cases the subject's cognitive abilities still play *some* role in explaining their success. For example, in typical Gettier cases the subject will employ their perceptual abilities, or their reasoning abilities. The same is true of the communicative credit theory: in order to properly deal with standard Loar cases it has to be maintained that the hearer's abilities are the most salient explanatory factor in their successful recovery of the speaker's intended message. After all, the audience's abilities still play some small role in these cases. For example, the audience's knowledge of the English language, and their general ability to infer what a speaker means by a use of 'that' on the basis of their contextual knowledge both play important roles in Loar's original case. This gives rise to two problems.

Firstly, in typical communicative exchanges the speaker's abilities will be just as important as the hearers in explaining the success of the communicative exchange. So, if understanding requires that the audience's abilities be the most salient explanatory factor, understanding will rarely be achieved.¹⁴ Secondly, there look to be cases of understanding in which some factor other

¹⁴ A parallel problem is raised by Lackey (2007) for the credit view of knowledge: in typical testimonial exchanges the speaker deserves as much or more credit for the audience's belief than the audience. There are many responses

than either the hearer or speaker's abilities is most salient in explaining the audience's success. For example, in German the most salient factor in explaining the audience's success is the lucky manner in which they formed their true belief about the meaning of 'Krankenschwester'. Yet, this is still a case of understanding.

Indeed, German seems to create a general problem for attempts to provide a unified response to the value problems for knowledge and understanding. The problem as it applies to the credit theory is as follows: In German the audience understands the speaker, but they are not in a position to know what the speaker has said. If the audience's abilities were sufficiently explanatorily central for them to deserve credit for their successful recovery of the speaker's intended meaning, then they would also be sufficiently explanatorily central for them to deserve credit for their accurate belief regarding the speaker's intended meaning. But if this was the case then they would be in a position to know what was said. Yet, they are not in a position to know what was said. So, the credit approach can't explain both the value of knowledge and the value of understanding.

This point seems to generalize. Any attempt to port over our favored response to the value problem for knowledge to the value problem for communication will have to contend with the fact that in **German** the subject understands but is not in a position to know what has been said. Suppose that we say that knowledge is valuable because it has some feature F. We would have to say that understanding is valuable due to its possessing F or some closely related feature F*. Since the subject in German is not in a position to know what has been said we will have to deny that their judgement that 'S said that p' has feature F. But since they do understand the utterance it looks like we will have to say that their entertaining of p does possess F, or some closely related feature F*. This suggests that it will be difficult to find parallel solutions to the value problems for knowledge and communication.

Hyska (MS) frames her challenge as follows: 'understanding bears a relationship to knowledge such that skepticism about the value of knowledge suggests a skepticism about the value of communication' (p 9). But the considerations raised above suggest that the challenge is more general. Even if we are not skeptics about the value of knowledge there are grounds to worry about the value of understanding.

In the final sections of this paper, I will explore two alternative ways of responding to the value problem for understanding. I will suggest some problems for the first approach (although these problems are by no means decisive), and I will tentatively endorse the second approach.

7 Representational Calibration

What is the main difference between **Loar** and **German**? Well, in both cases there is a clear sense in which the audience could easily have been led into error. And in neither case is the hearer's success an achievement (at least, not in the sense in which knowledge is often thought to be an achievement). Rather, the difference seems to be that in **Loar**, but not in **German**, the direct basis of the audience's interpretation could easily have led them to error.

to this objection, and some of them may carry over to the credit theory of understanding. So I do not want to place too much weight on this objection

In light of this, the following line of reasoning suggests itself: the value of understanding resides at least in part in the relationship between the direct basis of the audience's interpretation and the relevant utterance. But how could the value of understanding derive from this relationship? Here is a suggestion: When we communicate we aim to align our means of representation to those of our community. Successful understanding manifests the calibration of our interpretative mechanisms to our linguistic environment. That is, it manifests the alignment of our means of representation to our community.

This would explain the difference between **Loar** and **German**: In **Loar** the audience's successful interpretation of the speaker does not manifest the calibration of their interpretative mechanisms to their communicative environment. After all, given the direct basis of their interpretation they could easily have gone wrong. However, in **German** the subject's interpretative mechanisms are well attuned to their environment, and their successful recovery of the speaker's intended meaning manifested that fact.

This seems promising. However, it immediately raises a new question: Why would such calibration be among the aims of linguistic communication? Why would we not simply aim at, say, information transfer?

There are a few possible answers here. It could be, for example, that there is some special value in our means of representation being calibrated to linguistic practices that go back generations, and that are in some sense definitive of the communities of which we are members. A somewhat less nebulous answer can be found by examining an analogous question regarding epistemic evaluations.

Sinan Dogramaci (2012) observes that we have a practice of negatively appraising certain acts of belief formation even when they reliably track the truth. And we have a practice of praising certain acts of belief formation even when they fail to reliably track the truth. This is puzzling. After all, we aim to form true beliefs. If a method of belief formation reliably yields true beliefs for a subject, then why not hold it in high esteem?

Dogramaci's answer is that we use evaluations such as 'rational' or 'irrational' to influence others' behavior. When we credit someone with rationality we re-enforce their behavior. When we criticize someone as irrational, we do the opposite. In doing so, we pressure others to adopt the same patterns of reasoning that we employ. This in turn promotes the coordination of epistemic rules employed within a community. This coordination is desirable, according to Dogramaci, because it allows us to treat others as epistemic surrogates, and thus defer to them unproblematically.

Our question is similar to Dogramaci's. If an audience correctly grasps the proposition the speaker intends to communicate, then why would it matter how they got there? Why do we attribute understanding (a positive evaluation of their performance) only when the audience's success results from the calibration of their means of representation to their linguistic environment? Well, a natural response, following Dogramaci, is that in doing so we re-enforce their behavior. Our practice of positively evaluating instances of successful coordination specifically when they arise from the calibration of the audience's means of representation to their linguistic environment promotes the calibration of means of representation within the linguistic

community. And this, of course, is good for the community in general: it supports smooth successful communication within the community in at large.

This response appears promising. Unfortunately, this appearance is deceptive. We are trying to explain why it is better for the various conditions on understanding to be satisfied in a non-lucky way. The suggestion is that when we communicate, we aim for our coordination to manifest the calibration of our means of representation to our environment. And this is supposed to be a good thing to aim for because it promotes the calibration of representational devices within the community. However, if this is right, it is not clear why we would attribute understanding only when our coordination *manifests* the calibration of our means of representation to our environment. Why not attribute understanding whenever the following conjunctive condition is satisfied: 1) The hearer accurately recovers the speaker's intended proposition, and 2) the hearer's means of representation are calibrated to their linguistic environment? This would promote the coordination of means of representation just as well. However, it would also lead to attributions of understanding in cases like **Loar**. After all, **Loar** is the communicative equivalent of a Gettier case. It is a case in which the audience is both successful, and skillful. Moreover, their abilities are underwritten by the calibration of their means of representation to their linguistic environment. Jones employs his typically reliable abilities to recover the proposition intended by Smith. However, his success does not manifest his ability. Likewise, it doesn't manifest the calibration of his means of representation to his linguistic environment.

So, the least nebulous way of making sense of the representational calibration response quickly runs into problems. Perhaps these problems can be overcome, or perhaps there are other more promising ways of developing this strategy. But rather than exploring this possibility I'll close with what I consider to be a more promising alternative.

8 Intentional Coordination

So far, we have focused on analogies and disanalogies between communicative and epistemic luck. However, luck also plays a central role in the theory of intentional action: if an agent *s* with the intention of *ϕing*, but their *ϕing* was lucky, this will often render their action unintentional. Consider an example (C.f. Harman (1976)):

Unintentional Fulfilment: Tom intends to kill Pete. He starts driving to Pete's home with the intention of doing so, but on the way, he accidentally hits Pete with his car. Pete dies as a result.

Tom set off with the intention of killing Pete. Moreover, he did kill Pete. So, his intention was satisfied. However, he did not intentionally kill Pete. It was merely by luck that he ended up killing him.¹⁵

Communication is an intentional activity. Speakers have communicative intentions: intentions that their audience entertain certain propositions. Moreover, just like any other action, these intentions can be satisfied unintentionally. That is, a speaker's communicative intention can be satisfied (i.e. the hearer can entertain the intended proposition) without the speaker intentionally bringing it

¹⁵ Whilst I take this to be the standard interpretation of such cases there are exceptions. Some, such as Harman (1976) and Roth (2000) hold that the subject's intention is not actually satisfied in such cases.

about that their intention is satisfied (i.e. without the speaker intentionally bringing it about that the audience entertain the intended proposition). This is just what happens in cases like **Loar**. Smith intended that Jones entertain the proposition that a particular man is a stockbroker. Jones did entertain this proposition. Yet Smith's success is lucky in a way that seems to render it unintentional. Smith did not intentionally bring it about that Jones entertain this proposition. **Loar** seems directly analogous to **Unintentional Fulfilment** in this sense.

So, the following possibility presents itself: we could maintain that an audience understands an utterance only when the speaker intentionally brings it about that they entertain the intended proposition. What about cases like **German** in which there seems to be lucky understanding? Is the speaker's success there intentional? Well, consider the following case:

Fire: Yannick intends to start a fire. In order to start a fire, one must be in an oxygen rich environment. Unbeknownst to Yannick, he has entered an environment in which oxygen levels are typically very low, such that striking a match would not be sufficient to start a fire. However, luckily, and also unbeknownst to Yannick, some oxygen canisters have recently fallen out of a passing airplane and burst, leaking oxygen into his local environment and creating a small island in which the oxygen levels are very high. He lights his match and thereby starts a fire.

There is a clear sense in which Yannick's success is lucky. Nonetheless, his success was still intentional: he intentionally lit the match. What is distinctive about **Fire**? Well, in **Fire** the background conditions for success were established luckily. But once they were in place everything went as it should. Holding fixed the background conditions in **Fire** Yannick's success was not lucky. This sets **Fire** apart from both **Loar** and **Unintentional Fulfilment**. In both **Loar** and **Unintentional Fulfilment** luck intervenes in the process by which the intention is fulfilled. This seems to be inconsistent with intentional success.

Importantly, **German** is analogous in this respect to **Fire**, not to **Unintentional Fulfilment**: in **German** our protagonist is disposed to interpret 'krankenschwester' as 'nurse'. Their interpretative mechanisms are, in this sense, calibrated to their linguistic environment. The audience's interpretative mechanisms being well calibrated in this way is a background condition against which we typically form our communicative intentions. So, in **German**, like in **Fire**, it is merely a matter of luck that the background conditions for success are satisfied. But holding the background conditions in **German**, fixed it was not lucky that the speaker was successful. So, in **German**, unlike **Loar** the speaker intentionally brought it about that the audience entertain a particular proposition.

This is promising. Perhaps the value of understanding derives from the fact that understanding is a state intentionally brought about by a speaker. This would explain why it is bad for an audience to recover the speaker's intended proposition in a relevantly lucky way: when recovery occurs in a relevantly lucky way the speaker's success is unintentional. However, a new question immediately arises: why is it better for success to be intentional? In particular, why is it better for the *audience* if the *speaker* is intentionally successful rather than unintentionally successful?

A natural first thought here is that communication is a joint activity. The audience must do their part in interpreting the speaker. So, really, when they understand an utterance, it is not just the speaker's intentions that are satisfied: the audience's intentions (or, perhaps, their collective intentions) are satisfied too. Insofar as we care about our own successes being intentional, this would explain why it is better for the audience to understand the speaker, rather than satisfy all the other conditions on understanding minus the reliability condition.

I think there is a lot to be said for this approach, but I have some worries. Firstly, it is still not entirely clear why it is better to succeed at an activity intentionally rather than unintentionally (unintentional success does not warrant credit, but it is not clear that all intentional successes warrant credit either). Secondly, and more importantly, it is not clear that communication is always a jointly intentional activity. In many paradigmatic cases of utterance understanding the audience is passive. They merely hear an utterance and, without exercising any agency, represent the speaker as having asserted some proposition. For example, if I hear the utterances of a loud passenger on the train, I can understand their utterances despite having no interest in (and perhaps a desire to avoid) doing so. In these cases, it seems odd to describe the audience's interpretative success as intentional. And it is especially odd to describe the interaction between speaker and hearer as jointly intentional. Yet, understanding occurs here just as it does in active conversation.

The aforementioned problems are by no means decisive. However, I believe a more promising approach is available: *ceteris paribus*, when an agent intentionally ϕ s they bear a greater degree of responsibility for ϕ than when they unintentionally ϕ . Consider Tom and Pete in **Unintentional Fulfilment**: Tom was clearly responsible for killing Pete. However, had he killed Pete intentionally he would have borne a greater degree of responsibility for Pete's demise. We would consider him blameworthy to a higher degree. It is important to audiences that they be able to hold speakers responsible for the cognitive effects that utterances bring about (new beliefs, consideration of new questions, new plans etc.). Indeed, the greater the degree of responsibility a speaker bears for the cognitive change their utterance brings about, the better this is for the audience. After all, the greater the degree of responsibility the speaker bears for the audience's cognitive change (e.g. their new belief) the less responsibility the audience bears. So, it will generally be better for the audience if the speaker intentionally brings it about that they entertain some particular proposition than if they unintentionally bring it about that they entertain that same proposition, even if the speaker intends that they entertain that proposition in both cases. If this is right, then we have a neat solution to the value problem for understanding.

9 Conclusion

In this chapter I have done several things. I started by considering some reasons for and against the claim that utterance understanding must be reliable. There were two reasons for thinking understanding must be reliable. Firstly, understanding typically yields knowledge of what is said, and it typically supports the acquisition of testimonial knowledge. Secondly, there looked to be cases ('Loar cases') in which the unreliability of the audience's judgement seemed to preclude them from understanding an utterance. However, there were also reasons to doubt that understanding had to be reliable: there looked to be cases of lucky understanding. On closer investigation these cases didn't reveal that understanding could be lucky, but they did help us delimit the exact sense in which luck precludes understanding. This set us up for the second half of the chapter in which I discussed the question of why it is better to understand than to satisfy all

the other conditions of understanding in an unreliable way. I considered a number of solutions to this problem before tentatively endorsing the proposal that we understand only when a speaker intentionally brings it about that we entertain a particular content.

Acknowledgements

I would like to thank Megan Hyska, my colleagues on Prof. Robbie Williams's GROUNDS project at the University of Leeds, and the audience at the (Mis)communication and Context workshop at the University of Tartu for comments and discussion that greatly improved the quality of this paper and aided the development of the main thesis. This research was funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 818633)

References

- Broncano-Berrocal, F. 2018. Purifying impure virtue epistemology. *Philosophical Studies* 175(2). 385–410.
- Byrne, A., & Thau, M. 1996. In defense of the hybrid view. *Mind*, 105 (417), 139–149.
- Campbell, J. 1982. Knowledge and Understanding. *Philosophical Quarterly* 32 (126). 17-34.
- Davies, M. 1989. Tacit knowledge and subdoxastic states. In A. George (Ed.), *Reflections on Chomsky* (pp. 131–152). Oxford: Basil Blackwell.
- Dogramaci, S. 2012. Reverse Engineering Epistemic Evaluations. *Philosophy and Phenomenological Research* 85 (3). 513-530.
- Dummett, M. 1978. Frege's distinction between sense and reference. In his *Truth and other Enigmas*. London: Duckworth.
- Dummett, M. 1991. *The logical basis of metaphysics*. Cambridge, MA: Harvard University Press.
- Evans, G. 1982. *The varieties of reference*. Oxford: Clarendon Press.
- Fricker, E. 2003. Understanding and knowledge of what is said. In A. Barber (Ed.), *Epistemology of language* (pp. 325–366). Oxford: Oxford University Press.
- Goldman, A. 1986. *Epistemology and Cognition*. Cambridge: Harvard University Press.
- Graham, P. 2016. Against Actual World Reliabilism: Epistemically Correct Procedures, Reliably True Outcomes. In Angel Fernandez Vargas(ed.) *Performance Epistemology: Foundations and Applications*. Oxford, Oxford University Press.

- Graham, P. 2017. Normal Circumstances Reliabilism: Goldman on Reliability and Justified Belief. *Philosophical Topics* 45 (1). 33-61.
- Greco, J. 1999. Agent reliabilism. *Philosophical Perspectives* 33. 273–96. 25
- Greco, J. 2000. *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*. Cambridge: Cambridge University Press.
- Greco, J. 2003. Knowledge as Credit for True Belief. In DePaul, M., & Zagzebski, L. (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford: Oxford University Press. 111-134.
- Greco, J. 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.
- Grice, H. P. 1957. Meaning. *Philosophical Review* 66. 377-388.
- Harman, G. 1976. Practical Reasoning. *Review of Metaphysics* 29 (3). 431-463.
- Heck, R. K. 1995. A sense of communication. *Mind*,104(413). 79–106.
- Higginbotham, J. 1992. Truth and understanding. *Philosophical Studies* 65. 3–16.
- Hunter, D. 1998. Understanding and belief. *Philosophy and Phenomenological Research* 58 (3), 559–580.
- Hyska, M. Manuscript. Luck and the Value of Communication. Keiser, J. Forthcoming. Language Without Information Exchange. *Mind and Language*.
- Kelp, C. 2011. In defense of virtue epistemology. *Synthese* 179. 409–33.
- Kelp, C. 2016. Justified belief: knowledge first-style. *Philosophy and Phenomenological Research* 93. 79–100.
- Kelp, C. 2019. How to be a Reliabilist. *Philosophy and Phenomenological Research* 98 (2). 346-374.
- Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Lackey, J. 2008. What luck is not. *Australasian Journal of Philosophy* 86(2), 255–267.
- Leplin, J. 2007. In Defence of Reliabilism. *Philosophical Studies* 134. 31-42. Loar, B. 1976. The Semantics of Singular Terms. *Philosophical Studies* 30 (6). 353-377.
- Longworth, G. 2008. Linguistic understanding and knowledge. *Nous*, 42, 50–79.

Longworth, G. 2009. Some Models of Linguistic Understanding. *Baltic International Yearbook of Cognition, Logic, and Communication* Vol 5. 1-56.

26

Longworth, G. 2018. Understanding What is Said. *Synthese* 195. 815-834.

McDowell, J. 1994. Knowledge by hearsay. In B. K. Matilal A. Chakrabarti (Eds.), *Knowing from words: Western and Indian philosophical analyses of understanding and testimony* (pp. 195–224). Dordrecht: Kluwer. Reprinted in his 1998: 414–443.

Nozick, R. 1982. *Philosophical Explorations*. Cambridge: Harvard University Press.

Peet, A. 2018. Testimonial Knowledge without Knowledge of What is Said. *Pacific Philosophical Quarterly* 99 (1). 65-81.

Peet, A. 2019. Knowledge-Yielding Communication. *Philosophical Studies* 176. 3303-3327.

Peet, A. Forthcoming. Assertoric Content, Responsibility, and Metasemantics. *Mind and Language*.

Pettit, D. 2002. Why Knowledge is Unnecessary for Understanding Language. *Mind* 111 (3). 519–550.

Pollock, J. Forthcoming. Content Internalism and Testimonial Knowledge. *Inquiry* 1-22.

Pritchard, D. 2005. *Epistemic Luck*. Oxford: Oxford University Press.

Roth, A, S. 2000. The Self Referentiality of Intentions. *Philosophical Studies* 97.

11-52.

Smith, M. 2016. *Between Probability and Certainty: What Justifies Belief*. Oxford: Oxford University Press.

Sosa, E. 1991. *Knowledge in Perspective. Selected Essays in Epistemology*. Cambridge: Cambridge University Press.

Sosa, E. 1999. How to Defeat Opposition to Moore. *Philosophical Perspectives* 13. 137-149.

Sosa, E. 2010. How competence matters in epistemology. *Philosophical Perspectives*. 24. 465–475.

Sosa, E. 2011. *Knowing Full-Well*. Princeton/NJ: Princeton University Press.

Sosa, E. 2015. *Judgment and Agency*. Oxford: Oxford University Press.

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

Zagzebski, L. 2003. The Search for the Source of Epistemic Good. *Metaphilosophy* 34 (1/2). 12-28.

Zagzebski, L. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press.