

Assertion Compatibilism

Mona Simion

COGITO Epistemology Research Centre

University of Glasgow

Abstract

The vast majority of philosophers accept Assertion Incompatibilism: according to this view, given intuitive variability of proper assertion with practical stakes, non-shifty invariantism (NSI) is incompatible with a biconditional knowledge norm of assertion (KNA). There are also a few dissenting voices, however: some invariantists venture to explain the sensitivity data for proper assertion in a fashion that preserves both NSI and KNA (Assertion Compatibilism). In this paper, I argue that my preferred incarnation of Compatibilism fares better than the competition. According to the competition, shiftiness in proper assertability is to be explained via appealing to the pragmatics of language. According to the view I defend, what varies with practical considerations is the all-things-considered propriety of assertion: epistemic propriety and the epistemic standard at stake are invariant.

1. Introduction

Here are two attractive theses:

The Knowledge Norm of Assertion (KNA): One's assertion is epistemically permissible iff one knows that

p.

and

Non-Shifty Invariantism (NSI): The truth value of knowledge claims is insensitive to practical matters.

Now, theoretical considerations speak in favour of not quickly abandoning either of these two theses. First, friends and foes alike agree that NSI is the default epistemological position:¹ we need to be argued out of it. Second, if knowledge is both necessary and sufficient for epistemically permissible assertion,² we have a very straightforward and elegant way of explaining quite a few otherwise puzzling linguistic data, such as: the paradoxical soundingness of Moorean statements of the form ‘*p* but I don’t know that *p*’; the unassertability of lottery propositions; ‘How do you know?’ challenges; (intuitively permissible) criticisms such as ‘Why didn’t you say so? You knew it all along!’

Alas, popular as they might be, for the most part, the epistemological literature of the last decade takes KNA and NSI to be incompatible. The culprit is the intuitive sensitivity of permissible assertion to practical stakes. Roughly, here is the thought behind the incompatibility claim: Since permissible assertion seems to require more warrant in high stakes than in low stakes scenarios, we’ll have to choose between one of the following two options: We embrace KNA but then we’ll

¹ “Anyone who has even passing knowledge of analytic epistemology in the last fifty or so years knows that moderate invariantism is the orthodoxy. It is the view to beat” (Fantl and McGrath 2009, 37).

² Champions of KNA include DeRose (e.g. 2002), Hawthorne (e.g. 2004), Kelp (e.g. 2016), Simion (e.g. 2016, 2019, 2020), Turri (2011), Unger (1975), Williamson (1996, 2000).

have to adopt a view that takes knowledge/knowledge attribution to be sensitive to practical considerations. Alternatively, we stick to our non-shifty invariantist guns, but then we'll have to give up the knowledge norm of assertion. Let's call the view that KNA and NSI are incompatible 'Assertion Incompatibilism.'

While the vast majority of philosophers accept Assertion Incompatibilism, there are also a few dissenting voices: some invariantists venture to explain the sensitivity data for proper assertion in a fashion that preserves both NSI and KNA. Let us dub this position 'Assertion Compatibilism.' One notable compatibilist strategy on the market explains why proper assertion varies with stakes in terms of what is pragmatically conveyed by the assertion in question rather than in terms of what is, strictly speaking, said (e.g. Rysiew (2001), Brown (2006), Hazlett (2009), Kyriacou (2019), Unger (1975)). It is argued that neither knowledge (nor 'knowledge') nor proper assertion are sensitive to practical stakes: rather, the felt impropriety sometimes associated with knowledgeable assertions in high stakes is triggered by the generated implicatures. Let us call this view Pragmatic Compatibilism (henceforth also PC).

This paper argues that my preferred variety of Compatibilism – Functionalist Compatibilism – compares favourably to the PC competition. On Functionalist Compatibilism, what explains the intuition of shiftiness with stakes is the all-things-considered propriety rather than the epistemic propriety of assertion. Here is the game plan: Section #2 outlines the Incompatibilist challenge. In #3, I look at Pragmatic Compatibilism and argue that the view fails to generalize to knowledge-attribution-free assertions. Section #4 outlines my Functionalist Compatibilist response. In the

last section I conclude.

2. Contextualist Cases and the Generality Objection

Intuitively, proper assertion is epistemically more expensive when the stakes are higher in the sense that it appears to require more warrant. Let us dub this the Shiftiness Intuition.³ The Shiftiness Intuition is not exactly breaking news in epistemology.⁴ That said, it came into popularity when epistemic contextualists started using it in support of their view. Epistemic contextualism is a semantic thesis about attributions of knowledge: it holds that the meaning of ‘knows’ varies with context. It is typically supported by pairs of cases like the following ‘bank cases’ by Keith DeRose (1992: 913):

BANK CASE A. My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, ‘Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays.’ I reply, ‘No, I know it’ll be open. I was just there two weeks ago on Saturday. It’s

³ To my knowledge, the term was coined in (Fantl and McGrath 2012).

⁴ Austin (1979, 180), for instance, observes that, while in normal contexts the fact that your hat is in the hall seems to be good enough reason for me to say that you are in, when a lot hinges on it, I would be quite reticent to do the same.

open until noon.'

BANK CASE B. My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a very bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, 'Banks do change their hours. Do you know the bank will be open tomorrow?' Remaining as confident as I was before that the bank will be open then, still, I reply, 'Well, no, I don't know. I'd better go in and make sure.'

Let's assume that the bank will, in fact, be open on Saturday and so DeRose's corresponding belief is true in both CASE A and CASE B. While DeRose's attribution of knowledge that the bank is open tomorrow is intuitively permissible in CASE A, the same does not hold for CASE B.

Contextualism has a straightforward explanation of these data: the truth conditions of knowledge attributions vary with context with the result that, in CASE A, DeRose's attribution of knowledge is true, whereas in CASE B it is false.

Now, here is one important difficulty for epistemic contextualism, which DeRose himself recognizes and dubs The Generality Objection: the shiftiness intuition does not only arise for knowledge attributions. Rather, it

is a much more general phenomenon. For instance, in CASE A, it is not only fine for DeRose to assert that he knows that the bank will be open tomorrow, it would also be fine for him to assert that the bank will be open tomorrow. In contrast, in CASE B, it would not only seem wrong for DeRose to assert that he knows that the bank will be open tomorrow, but also to assert that the bank will be open tomorrow (DeRose 2002, 177). To see this more clearly, consider the following cases:

ASPIRIN₁. You remember having bought aspirin last month. You are heading together with your sister towards your place for dinner, and she lets you know she has a minor headache. Your sister asks you: *‘Do you have aspirin at home, or should we go to the pharmacy?’* You flat out assert: *‘Don’t worry, I have aspirin at home, I remember having bought some’*.

ASPIRIN₂. You remember having bought aspirin last month. Your sister’s two years old baby is having a fever, and needs an aspirin as soon as possible. Your sister asks you: *‘Do you have aspirin at home, or should we go to the pharmacy?’*. You give the matter a bit of thought, and answer: *‘Well, let’s drop by the pharmacy, just in case’*.

Intuitively, in ASPIRIN₁, but not in ASPIRIN₂, it is appropriate to flat out assert that you have aspirin at home. What these cases suggest is that the shiftiness intuition is not limited to cases of knowledge attributions. As a result, the worry is that epistemic contextualism, in virtue of being a view about the semantics of ‘know,’ won’t give us the right explanation of the shiftiness intuition after all.

In order to address this worry, contextualists have

appealed to the knowledge norm of assertion (KNA).⁵ Roughly, here is the thought: according to KNA, one's assertion that p is epistemically proper if and only if one knows that p . If that is the case, however, it follows that the standards for knowledge go hand in hand with the standards for proper assertion. But now it looks as though epistemic contextualism is well positioned to avoid the generality objection. After all, the objection suggests that the standards for proper assertion vary in general (that is, not just in cases featuring knowledge attributions). But if KNA and contextualism are both true, this is just what we'd expect anyway. In fact, KNA and contextualism in conjunction promise to offer an appealing explanation of the relevant data. To see this, consider first the variation of CASE A in which DeRose asserts that the bank is open. Suppose, plausibly enough, that the bar for 'knowledge' is low here, low enough that DeRose 'knows' what he asserts and so his assertion comes out proper. In contrast, suppose DeRose were to assert that the bank is open in a variation of CASE B. Suppose, again plausibly, that the bar for 'knowledge' is high here, high enough that DeRose doesn't 'know' what he asserts and so his assertion would come out improper. In this way, KNA and contextualism together can offer an appealing solution to the generality problem.

In fact, there is reason to think that DeRose can now turn the tables on the generality objector and argue that KNA demands an account of knowledge according to which whether one knows is sensitive to practical stakes:

What of the advocate of the knowledge account of

⁵7 The locus classicus for the defence of the necessity claim involved in KNA is Williamson (2000). For support for the sufficiency claim, see Simion (2016a).

assertion who does not accept contextualism? Such a character is in serious trouble. Given invariantism about knowledge, the knowledge account of assertion is an untenable attempt to rest a madly swaying distinction upon a stubbornly fixed foundation. [...] The knowledge account of assertion demands a contextualist account of knowledge and is simply incredible without it (2002, 182).

If DeRose is right, KNA is incompatible with Non-Shifty Invariantism: since the standards for knowledge and permissible assertion co-vary, and since the latter are shifty with stakes, it seems to follow that the standards for knowledge will be shifty with stakes as well.

3. Pragmatic Compatibilism, Modest and Sceptical

According to Pragmatic Compatibilists, the source of variability in contextualist cases pertains to what is pragmatically conveyed by the assertion in question rather than by what is, strictly speaking, said (e.g. Brown 2006, Hazlett 2009, Kyriacou 2019, Rysiew 2001, Unger 1975).

Non-shifty invariantism comes in two different flavours: sceptical and moderate. The sceptical invariantist (e.g. Kyriacou 2019, Schaffer 2004, Unger 1975) claims that the semantic value of the word 'know' is such that all or nearly all ordinary positive knowledge ascriptions of the form 'S knows that p ' are false. The moderate invariantist claims that the semantic value of 'know' is such that many of the positive knowledge

ascriptions that we make in daily life are true. Both views are Non-Shifty varieties of invariantism.

Correspondingly, Pragmatic Compatibilism comes in two different flavours as well. On Moderate Pragmatic Compatibilism (Brown 2006, Hazlett 2009, Rysiew 2001, henceforth also MPC), in both CASE A and CASE B, DeRose does, in fact, know that the bank will be open on Saturday. However, given the presumption of relevance and informativeness, in CASE B, saying that he knows pragmatically conveys that he is able to eliminate all contextually salient error possibilities – such as, for instance, the possibility that the bank changed its hours in the last two weeks. But of course, this is just an error possibility that DeRose cannot eliminate. In this way, then, DeRose’s self-attribution of knowledge would carry a false implicature. And it is just this false implicature that, according to Pragmatic Compatibilists, explains why DeRose’s self-attribution of knowledge would seem wrong.

On this view, Non-shifty Invariantism is perfectly compatible with KNA: DeRose knows in both CASE A and CASE B and is, therefore, in a good enough epistemic position to assert that the bank will be open.

Sceptical Invariantism (Kyriacou 2019, Schaffer 2004, Unger 1984,) traditionally has a somewhat harder time accommodating the shiftiness intuition without giving up on KNA. According to this view, speakers ascribing knowledge speak non-literally, i.e. the speaker meaning does not coincide with the meaning of the sentence uttered. What varies with context, according to these philosophers, is not what the speaker literally says, i.e. the truth-conditional content of the sentence used, but what she means by her use of that sentence. According to Sceptical Invariantism, by uttering ‘x knows that *p*,’ what

the speaker says is ‘x can rule out every conceptually possible alternative to *p*,’ while what she means is ‘x can rule out every *relevant* alternative to *p*.’ As such, truth conditions come apart from assertability conditions; KNA is thus not easily vindicated by this view. Recently, though, Christos Kyriacou (2019) argues for a Sceptical Pragmatic Compatibilism (henceforth SPC): on this deflationary view, although strictly speaking we do not know most of the things we take ourselves to know, KNA concerns a weaker, every-day notion of knowledge – henceforth weak-knowledge that is useful to have, and that we employ in our everyday affairs. As such, on this view, it is this weaker notion that is at stake in KNA, and it is the very same notion that is responsible, just like in the case of Modest Pragmatic Compatibilism, for the implicatures allegedly to blame for generating the shiftiness intuition.

A lot of ink has been spilled on whether the Pragmatic Compatibilist account of the bank cases is ultimately satisfactory for accounting of the shiftiness of proper *knowledge attributions*.⁶ I will not address this issue here. In contrast, what I will focus on is whether Pragmatic Compatibilism can explain the shiftiness intuition *in its full generality*. After all, we have seen that this intuition is not restricted to knowledge attributions, but rather it also arises in cases of assertions of propositions that do not feature epistemic concepts. By way of illustration, recall once more the variation of CASE B in which were DeRose to assert simply that the bank is open tomorrow, his assertion would be improper. Note that Pragmatic Compatibilists will have a hard time extending their account of the original CASE B to this variation. The thought that the knowledge attribution, in

⁶ See e.g. (DeRose 2009).

conjunction with contextually salient error possibilities, generates an implicature to the effect that DeRose can eliminate the salient error possibilities has little traction here. After all, there is no knowledge attribution being made.

To the best of my knowledge, Hazlett (2009, 31) is the only champion of Pragmatic Compatibilism that has addressed this issue in print, on behalf of the Moderate variety of the view, so in what follows I will first examine his solution in detail. Further on, though, I will also look at one more possible response on behalf of MPC, as well as a response on behalf of SPC – the sceptical version of the view.

Here is Hazlett: say that, in the relevant variation of CASE B, DeRose were to assert that the bank is open tomorrow. Now, by the Gricean maxim of Quality (alternatively, if knowledge is the norm of assertion), this generates the implicature that DeRose knows that the bank is open tomorrow. According to Hazlett, this, in turn, will imply that the speaker can eliminate all the contextually salient error possibilities. For instance, in CASE B, were DeRose to assert that the bank will be open tomorrow, his assertion would generate the implicature that he knows this, which, in turn, would generate the implicature that he can eliminate the contextually salient error possibility that the bank changed its hours in the last two weeks. Since DeRose can do no such thing, the assertion does generate a false implicature, i.e. in virtue of generating an implicature that itself has a false implicature.

3.1. Shiftiness Without Error Possibilities

One immediate worry with this line is that it is easy to come up with cases that are just like the variation of CASE B we have been considering in which there are no contextually salient error possibilities being tabled. The ASPIRIN cases we have been looking at are precisely like that. In these cases, in contrast to the Bank Cases, no error possibilities are raised; still, the Shifty Intuition survives.

Perhaps the Pragmatic Compatibilist could respond that the set of contextually salient error possibilities need not be made explicit. It will do if there is an implicit set of such possibilities. For instance, in ASPIRIN₂, the Pragmatic Compatibilist could perhaps claim that the possibility that you misremember having bought aspirin becomes implicitly salient.

The trouble with “going implicit” is that the solution threatens to be too easy. For this move to enjoy any plausibility whatsoever, we will need an independently plausible account of how the set of implicitly salient error possibilities is generated. Given that what we are dealing with is a context-sensitive phenomenon, it is quite plausible that it is facts about the participants to the conversation that generate the set, such as their interests, what they are attending to, etc. It is hard to deny that these facts vary from one speaker to the next, which begs the question as to just how a single set is settled on. One obvious proposal is that it is the lowest common denominator, as it were, i.e. the intersection of the set of error possibilities that each speaker would generate individually. However, this proposal is tricky because it is just not clear why all the error possibilities the Pragmatic Compatibilist needs will always be generated. For instance, why should we think that, in ASPIRIN₂, the possibility of misremembering

should always be generated as implicitly salient at this context? After all, since your sister has no idea when and if you bought aspirin – for all she knows, it may well be that you did so that very morning - , considerations pertaining to the reliability of your memory might not be particularly salient for her. Note that if it is possible that the possibility of misremembering is not part of the common denominator set, it looks as though the Pragmatic Compatibilist account of the data is in trouble.

Another obvious proposal is that it is the highest common denominator, as it were, i.e. the union of the sets of error possibilities that each speaker would generate individually. This proposal is also tricky because it is not clear why there couldn't be a participant to a low-stakes conversation whose interests etc. unduly raise the bar for 'knowledge'. But given all this, the Pragmatic Compatibilist who wants to solve the problem posed by cases like ASPIRIN₂ shoulders a substantive explanatory burden. Unless this burden has been discharged, we have every reason to remain sceptical about the Pragmatic Compatibilist line here.

3.2 Shiftiness Without Stakes

Cases like ASPIRIN₂ suggest that contextually salient error possibilities are not necessary for triggering the unassertability intuition. Even if it turns out that this problem can be circumvented, perhaps along the lines just outlined, the Pragmatic Compatibilist faces a further problem: they don't seem to be sufficient for doing the job either. That is, it looks as though, at least in some cases, tabling error possibilities will not do the trick, at least not in the absence of accompanying high stakes (see also

Brown 2006 for an acknowledgement of this). If that is the case, however, it starts looking as though proper assertability does not have all that much to do with salient error possibilities to begin with. Rather, it seems more and more plausible that the only thing that matters are variations in stakes. Thus consider the following variation of ASPIRIN₁:

ASPIRIN₁*. You remember having bought aspirin last month. You are heading together with your sister towards your place for dinner, and she lets you know she has a minor headache. You flat out assert: *'Don't worry, I have aspirin at home, I remember having bought some'*. Your sister asks: "Maybe you're mistaken; you do sometimes misremember things," she remarks. "I *have* aspirin at home," I reply. "All right," she says.

I take the intuition here to be that you and your sister are right to put the matter to rest. If that is true, however, tabled error possibilities by themselves – that is, in absence of raised stakes – won't seem to do the work Pragmatic Compatibilists takes them to do. This, of course, should hardly come as a surprise: just because sceptical worries are 'in the air' while I write an epistemology paper, for instance, it hardly follows I do not know and cannot assert that there is a computer in front of me. More is needed to trigger unassertability.

3.3 Implicature-Generating Implicatures and Sceptical Alternatives

We still have not reached the end of the obstacles the Pragmatic Compatibilist will have to overcome. To see

this, note that she is wheeling in some pretty heavy theoretical machinery, i.e. the idea of an implicature-generating implicature: for her account to work, the implicature that the speaker knows what she says (generated by KNA) in itself must be able to generate the further implicature that she can eliminate all contextually salient error possibilities.

Again, there is an explanatory burden to be discharged. To begin with, one might wonder whether there is such a thing as implicatures generating implicatures in the first place. Moreover, even if it turns out that the phenomenon exists, it would be good to know whether what we are dealing with here is a fully general phenomenon, or whether it is restricted to particular types of implicatures. And then there is the issue of just how implicatures generate implicatures. And finally, the account had better work, at least for the type of implicature that the Pragmatic Compatibilist wants to invoke.. In this way, the burden on the Pragmatic Compatibilist shoulders gets weightier and weightier.

Here is a more direct problem for the Pragmatic Compatibilist. It is widely agreed that implicatures must be calculable. And, of course, whether a given implicature is calculable will depend on the cognitive capacities of the agents who are facing the calculation task. But given that this is so, it is surely possible for there to be communities of speakers that are only able to generate simple implicatures, and not implicatures of implicatures. It is also surely possible for some such communities (i) to operate the concept of knowledge and have a word expressing it in their language and (ii) to have variations in stakes. If so, we can generate analogues of DeRose's bank cases for members of these communities. Note that the Pragmatic Compatibilist account will work just fine

for the analogues of DeRose's original bank cases. After all, all the Pragmatic Compatibilist needs here is a simple implicature. However, the envisaged account for the generalized cases like ASPIRIN₂ will just not work. After all, according to the Pragmatic Compatibilist line, what's going on here is an implicature of an implicature. Since members of the community cannot calculate implicatures of implicatures and since implicatures must be calculable, it follows that there is no such thing as an implicature of an implicature for this community. By the same token, the Pragmatic Compatibilist account is bound to fail here.

One could think that Sceptical Pragmatic Compatibilism will do better here, since they can have their cake and eat it too: it is open to these philosophers to argue that weak-knowledge, but not knowledge proper, is pragmatically encroached upon (i.e., a practically shifty relevant alternatives view of weak-knowledge, but not of knowledge, is correct). In turn, this view could be plausibly motivated by the thought that weak-knowledge is an entity that we only employ for practical purposes anyway. The problem with a view like this, however, is that it's not clear that it leaves room for any philosophically substantive – rather than merely terminological – debate between Sceptical Pragmatic Compatibilism and shifty incompatibilist views of knowledge, such as Pragmatic Encroachment. After all, it looks as though what Pragmatic Encroachment calls 'knowledge', the champion of skeptical pragmatic compatibilism calls 'weak-knowledge,' (in that both of these phenomena are pragmatically encroached). Also, the 'real knowledge' by SPC's lights looks suspiciously similar to what everybody else in the literature calls 'epistemic certainty.' If all this is so, the debate between Compatibilism and Incompatibilism becomes merely

terminological – which, in turn, I take it, will not be a satisfactory result for either side.

3.4. Shiftiness Without Error Possibilities

Last, I would want to look at an alternative possible reply on behalf of MPC that does not rely on the problematic notion of implicature-generating-implicature,⁷ and which, one might think, is closer to the spirit of MPC to begin with.

What seems key for people like Brown (2006) and Rysiew (2001) is that in the relevant cases (e.g., DeRose's Case B) what's practically important is that the subject be in a very good/strong epistemic position vis-à-vis the proposition in question. Due to considerations of relevance (it's alleged), a positive claim to knowledge in Case B would then 'pragmatically impart' that *S*'s epistemic position is good enough *relative to those heightened standards*, and so that *S* can rule out any salient alternatives. But this doesn't clearly mean that some such alternatives must be mentioned in advance for the former proposition, which is plausibly false, to be communicated. Why would it be improper for DeRose to assert, in B, simply that the bank is open tomorrow? Because he'd be representing himself (as per the maxim of Quality) as having 'adequate evidence', where such adequacy is (as per the maxim of Relevance) is context-, and so stakes-, sensitive. Given this, in asserting that *p*, DeRose would be pragmatically communicating that his epistemic position vis-à-vis *p* is much stronger than it plausibly is.

The worry for this reply is that the notion of 'contextually-sufficient' adequate evidence is itself typed: it can be sufficient e.g. prudentially – i.e., for practical purposes - morally – i.e., for moral purposes - , or epistemically – for

⁷ I would like to thank very much for this suggestion to an anonymous referee who reviewed my manuscript 'Shifty Speech and Independent Thought' for Oxford University Press.

epistemic purposes. What the champion of MPC needs, and is missing, is an argument that it is the former and *not the latter* sense of enough that is at stake here - on pain of either collapsing into pragmatic encroachment (if the former obtains), or begging the question against it (if no argument is given).

Furthermore, as I am about to argue, once we recognize and have an argument in favour of prudential normativity being at work in these cases, we no longer need to rely on the pragmatics of language to do the work in explaining the shiftiness intuition.

4. Functionalist Compatibilism

In previous work, (e.g. Simion 2018, 2019, 2020), I have defended a Compatibilist position that I dub ‘Functionalist Compatibilism.’ According to this account, once we distinguish between assertion’s epistemic function on one hand, and its prudential function on the other hand, the incompatibility worry sourced in the Shiftiness Intuition dissipates: KNA and NSI (in both of its incarnations) are perfectly compatible.

4.1 The Functionalist Framework

Here is the view, in a nutshell: The main epistemic function of assertion is to generate testimonial knowledge in hearers. In the vast majority of the cases, assertion will generate testimonial knowledge that *p* if and only if the speaker knows that *p*.⁸ Thus, in virtue of the function of assertion of generating knowledge in hearers, one is in a good enough epistemic position to make an epistemically proper assertion that *p* if and only if one knows that *p*. In

⁸ See Lackey (2008) for exceptions.

turn, whether one knows (or ‘knows’) that p is insensitive to practical matters. Crucially, on this account, what KNA claims is mere epistemic propriety. The way to ascertain whether the requirements at work in one case or another are genuinely epistemic requirements is by looking at the function that is plausibly being served. In traits, artifacts and actions alike, only functions of a type T generate norms of type T ; there is such a thing as a biologically properly functioning heart – associated with its biologic function of pumping blood –, a practically properly functioning can opener – borne out by its practical function of opening cans –, and an aesthetically proper way to engage in dancing waltz –, aimed at serving the aesthetic function of exemplifying beauty. We should expect, then, that if there is such a thing as epistemic propriety of assertion, it will somehow be associated with the epistemic function it serves.

Now, what epistemic goods is assertion meant to deliver? The plausible answer is that, characteristically, assertions will aim at generating testimonial knowledge in the audience.⁹ Due to our physical and cognitive limitations, a lot of the knowledge we have is testimonial; thus, assertion is one of our main epistemic vehicles.

Functional types in general come with norms of proper functioning. More specifically, these norms specify conditions that will ensure that the relevant function is served reliably enough. Note that Assertion will, with very few exceptions, generate knowledge in the hearer if and only if the speaker also knows; it should be pretty clear then that being knowledgeable is the most reliable way for assertion to fulfil its epistemic function. By the same

⁹ For arguments to this effect see, e.g., (Simion 2018, 2020), (Goldberg 2015), (Kelp 2016), (Reynolds 2002), (Turri In press)). See Graham (2010) for a defence of a true belief generation function of assertion.

token, there is reason to think that assertion is governed by a knowledge norm.

Now, importantly, notice that a given trait/artifact/act can have several functions simultaneously, even several functions of different types; take, for instance, the functions served by food for humans. One important such function will surely be a biological one, a nourishment function. Plausibly enough, though, on top of this, food also serves an aesthetic function, that of generating pleasant gustatory experiences. Now, normally, the aesthetic function complements the nourishing function. It serves survival by increasing the probability of us ingesting nourishing substances. This need not be the case, however; there can be situations where the two functions come in conflict, at which point the more stringent requirement will take precedence. In other words, when there is a conflict between the normative requirements associated with two functions, one requirement may override the other and dictate what's the all-things-considered good to observe. By way of illustration, think of a case where I am on a desert island and all I can eat in order to stay alive are my boots; surely, against my aesthetic well-being, that is exactly what I should do, all-things-considered.

To repeat, the thought here is that we get overriding from conflicts between the normative requirements associated with functions. We can now take this idea and put it to use for our epistemological purposes. To see how, note first that it is highly plausible that the epistemic function is merely one of the many functions served by assertion. One other very important function of assertion, as with action in general, will be a prudential one, serving our practical ends. Again, just like in the case of food, the epistemic function will, in most

cases, complement this prudential function. Generating testimonial knowledge in one's hearer with regard to an imminent threat, or about the whereabouts of resources are paradigm cases. However, again, this need not be the case. The two functions can also come into conflict. For instance, even if one knows that one's boss is bald, it may not be polite, prudent, or relevant to point this out to him ((Brown 2010, 550): here, the prudential function comes in conflict with the epistemic one. What's more, it's also plausible that the prudential function overrides the epistemic. The result is that although it is epistemically permissible to assert that one's boss is bald to his face, it is not all things considered permissible for one to do so.

One question that might arise at this stage is:¹⁰ if the epistemic function of assertion – that of generating testimonial knowledge – is its main function – which, plausibly, it is – why is it that norms generated by prudential normativity – so by a presumably secondary, prudential function of assertion – override the epistemic ones? Shouldn't norms generated by its main function be stronger?

Two things about this. First, the short answer is 'no.' In general, norms generated by the main function of a trait or artefact etc. need not override norms generated by secondary functions. Indeed, prudential and moral norms will plausibly have a high disposition to dictate all-things-considered-permissibility in the normative landscape, whether they are generated by main or secondary functions. Take, for instance, your washing machine: analytically, I take it, its main function is to wash laundry. If, however, laundry washing is dangerous to your life for some reason, the all-things-considered normative picture will be such that it is impermissible for

¹⁰ Thanks a lot to James Beebe for pressing me on this.

your washing machine to wash laundry.

Second, compatibly, notice that a washing machine that fails at washing laundry is, importantly, a bad token of its type, in a way in which a washing machine that fails to save our life is not. That is because, even though main functions need not generate all-things-considered prescriptive demands – i.e. concerning what ought-to-be-done – they do generate all-things-considered evaluative demands – i.e. concerning ought-to-bes. Ought-to-bes correspond to attributive goodness (a la Geach 1956): a washing machine that fails to wash but saves your life thereby is bad attributively, i.e. as a token of its type: it fails to be how it ought-to-be, although it does what it ought to do.

4.2 Modest and Sceptical Functionalist Compatibilism

The phenomenon of conflict and overriding is precisely what explains, in my view, the unassertability intuition in high stakes scenarios, no matter what flavour of Non-Shifty Invariantism one prefers – i.e., be it modes or sceptical. Since I myself have a mild preference for the former, I will start with the modest variety thereof. In a nutshell, the account predicts that, just like in the case of the washing machine, in the case of assertion, when the epistemic norm is overridden by practical considerations as in the Bank Cases, what one needs to do is not assert, even if one knows. Compatibly, though, if one were to (all-things-considered impermissibly) assert, one's assertion would be attributively good, i.e. a good token of its type (see also Kelp (2016), Simion, Kelp and Ghijsen (2016)).

On a modest invariantist variety of my Functionalism, then, in ASPIRIN₂, strictly epistemically

speaking, if you have memorial knowledge that you have aspirin at home, you are permitted to assert. You are in a position to make an assertion that can fulfil its epistemic function of generating testimonial knowledge in your hearer. However, prudential constraints concerning your nephew's health come into conflict with these epistemic requirements. They override the epistemic constraint and drive the degree of warrant needed for all-things-considered permissible assertion up to a point the speaker's epistemic support does not reach. As a result, it will be all-things-considered impermissible for the speaker to assert in such a case, which explains why her assertion would seem improper to us.

Similarly, DeRose does know that the bank will be open on Saturday in Bank Case B if he knows in Bank Case A. It is, though, prudentially and all-things-considered inappropriate for him to say so.

What about a sceptical variety of Functionalist Compatibilism? How would such a view go? I will start, once again, with my preferred variety of this view: on this account, standards for knowledge are very high – although not non-achievable. We can know from super-safe processes like deduction or perception of nearby large dry goods, but that's about it: knowledge from flimsy processes, such as long-term memory or testimony, is not easy to get. Notice that the view is not as far fetched as sceptical views are usually taken to be – in that it is much more in line with folk intuitions about knowledge: often ordinary folk are reticent to ascribe knowledge based on mere memory or testimony (see e.g. Lackey 2011 and Simion 2016 for discussion). To see it for testimony, think of a modified version of Lackey's (2008) case of the 'Chicago Visitor': upon arriving at the Chicago station, Morris gets off the train and asks a passer-by for

directions to the Sears Tower. The passer-by, knowledgeable about the geography of Chicago, tells Morris that the Sears Tower is to the left. Lackey's reported intuition (and mine) is that Morris thereby comes to know that the Sears Tower is to the left. The sceptical invariantist - and many undergraduate students, for that matter - will deny this claim. One way to see the sceptical point here is to imagine that, while on his way to the Sears Tower, a tourist mistakenly takes Morris to be a local and asks him which way to go to the Sears Tower. Many have the intuition that Morris should not outright assert 'To the left!,' but rather hedge along the lines of 'I'm not from around here myself, but a local told me it's to the left'. This intuition is water to the sceptical invariantist's mill.

On this view, De Rose does not know that the bank will be open on Saturday in either of the Bank Cases; After all, he's basing his corresponding assertion on a non-knowledge conferring memory from two weeks ago. Epistemically, then, he should not assert. That being said, practical considerations override the epistemic norm in both cases: in both cases, he is under practical pressure from his wife's question and the proximity of the bank to decide one way or another - to stop at the bank or not. In the first case, it's prudentially best if he does not stop - which explains the prudential permissibility of the corresponding assertion. In the second case, it is prudentially best if he stops - which is why it is prudentially impermissible to assert outright that the bank will be open tomorrow.

The last option that needs to be discussed is how Functional Compatibilism will deal with a strong version of sceptical invariantism, according to which we pretty much have no knowledge at all. Importantly for

our purposes here, note that the super-sceptical variety of this view can also afford the above explanation of the intuitive shiftiness in the bank cases, in terms of the prudential normativity at work in the contextualist cases. The extra problem, though, with this type of view, is that having a knowledge norm of assertion does not make sense to begin with on super-scepticism, since there is no such thing as knowledge to be had and to base one's assertions on. My Functionalism cannot solve this problem; neither is it, though, a problem that it is called to solve. The framework put forth in this paper merely concerns itself with rendering KNA and non-shifty invariantism compatible *in the presence of the shiftiness intuition*. What this latter problem makes clear, however, is that super-sceptical varieties of invariantism are not compatible with KNA on different grounds, independent of the shiftiness intuition (see also Hawthorne 2004). Solving this problem falls beyond the scope and the ambition of this paper.

5. Conclusion

This paper has defended Assertion Compatibilism: a biconditional norm of assertion need not imply a view about knowledge (or 'knowledge') that takes it to be sensitive to practical matters. I have argued here that there are problems with the extant compatibilist views on the market. Pragmatic Compatibilism – be it of the modest or sceptical variety –, attempting to explain the Shiftiness Intuition in Gricean terms, was shown to face the generality objection: that is, it fails to adequately account for variations in assertability in cases in which no epistemic concepts are deployed and no error possibilities

are on the table. I have argued that my favourite compatibilist view, Functional Compatibilism, compares favourably with the competition in that it succeeds in its compatibilist ambition. This view makes use of normative overriding for explaining the intuitive variability of proper assertion with stakes. What varies with practical matters is all-things-considered propriety: epistemic propriety, together with the epistemic standard at stake, are invariant.

References

- Brown, J. (2005). Williamson on Luminosity and Contextualism. *The Philosophical Quarterly*, Vol. 55, No. 219, 04.2005, p. 319-327.
- Capes, J (2012). Blameworthiness without Wrongdoing. *Pacific Philosophical Quarterly* 93 (3): 417-437
- DeRose, K. (2002). Assertion, Knowledge, and Context. *Philosophical Review* 111 (2), 167- 203.
- DeRose, K. (2009). *The Case for Contextualism: Knowledge, Skepticism, and Context: Vol. 1*, Oxford: Clarendon Press.
- Douven, I. (2006). Assertion, Knowledge, and Rational Credibility. *Philosophical Review*. 115 (4), 449-86.
- Fantl, J. and McGrath, M. (2012) Arguing for Shifty Epistemology. In Brown & Gerken (eds.), *Knowledge Ascriptions*. Oxford University Press 55–74.
- Graham P.J. (2010). Testimonial entitlement and the Function of Comprehension. In Duncan Pritchard, Alan Millar & Adrian Haddock (eds.), *Social Epistemology*.

- Oxford University Press 148--174 (2010)
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hazlett, A. (2009). Knowledge and Conversation. *Philosophy and Phenomenological Research*, 78(3): 591-620.
- Kelp, C. (2016). *Assertion: A Function First Account*. *Nous*, Online First.
- Kyriacou, C. (2019). Assertion and Practical Reasoning, Fallibilism and Pragmatic Skepticism. *Acta Analytica*, Online First.
- Kluwe, R.H. (1982). Cognitive Knowledge and Executive Control: Metacognition. In Griffin, D.R. (ed.) *Animal Mind – Human Mind*. Springer: 201-224.
- Lackey, J. (2007). Norms of Assertion. *Noûs* 41: 594-626.
- Lackey, J. (2008). *Learning from Words*, Oxford: Oxford University Press.
- Lackey, J. (2011). Assertion and Isolated Second-Hand Knowledge. In Brown, J. and Cappelen, H. 'Assertion', Oxford: Oxford University Press.
- Lehrer, K. and Cohen, S. (1983). Justification, Truth, and Knowledge. *Synthese*, 55(2): 191–207.
- Schultz, M. (Forthcoming). Decisions and Higher-Order Knowledge. *Nous*.
- Forthcoming. A Social Epistemology of Assertion (with C. Kelp). *Oxford Handbook of Social Epistemology*, Lackey, J. and McGlynn, A. (eds.). Oxford: Oxford University Press.
- Simion, M. (2020). Shifty Speech and Independent Thought. Manuscript.
- Simion, M. (2019). Saying and Believing: The Norm Commonality Assumption. *Philosophical Studies* Volume 176, Issue 8, pp 1951–1966. Online

- Simion, M. (2018). Assertion: The Context Shiftiness Dilemma. *Mind & Language*, 34:503–517. Online
- Simion, M. (2016). Assertion: Knowledge is Enough. *Synthese*, 193(10), 3041-3056.
- Simion, M., Kelp, C. and Ghijssen, H. (2016). Norms of Belief. *Philosophical Issues, A Supplement to Nous: Epistemology and Mind Issue* (Jack Lyons and Chris Kelp, eds.). vol. 26/1: 375-392.
- Rysiew, P. (2001). The Context-Sensitivity of Knowledge Attributions. *Noûs*, 35(4): 477–514.
- Schaffer, J. (2004). Skepticism, Contextualism, and Discrimination *Philosophy and Phenomenological Research* Vol. 69, No. 1 (Jul., 2004), pp. 138-155
- Turri, J. (2010). Epistemic Invariantism and Speech Act Contextualism. *Philosophical Review*, 119: 77–95.
- Turri, J. (2013). Knowledge Guaranteed. *Noûs* 47: 602-612.
- Unger, P. (1975). *Ignorance: A Case for Scepticism*. Clarendon, Oxford.
- Weiner, M. (2005). Must We Know What We Say?. *Philosophical Review* 114 (2), 227-51.
- Williamson, T. (1996) Knowing and Asserting, *Philosophical Review* 105 (4), 489-523.
- Williamson, T. (2000). *Knowledge and Its limits*. Oxford: Oxford University Press.