Kelp, C. (2018) Assertion: a function first account. *Noûs*,

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

http://eprints.gla.ac.uk/140977

Deposited on: 16 May 2018

Assertion: A Function First Account

Christoph Kelp*

Abstract

This paper aims to develop a novel account of the normativity of assertion. Its core thesis is that assertion has an etiological epistemic function, *viz.* to generate knowledge in hearers. In conjunction with a general account of etiological functions and their normative import, it is argued that an assertion is epistemically good if and only if it has the disposition to generate knowledge in hearers. In addition, reason is provided to believe that it makes sense to regulate the practice of assertion by a speaker rule—and, more specifically, by a knowledge rule—as so regulating assertion contributes to ensuring that assertion fulfils its etiological function reliably.

Introduction

A question that has been at the forefront of recent epistemology concerns the normativity of assertion. According to a popular thesis, assertion is governed by the following epistemic rule:

THE C RULE OF ASSERTION. One must (epistemically): assert p only if p has C.

[Williamson 2000: 241]

It is commonly assumed that THE C RULE OF ASSERTION is the only epistemic rule governing assertion. In what follows, I will be working on this assumption as well.

In addition, we evaluate assertions. For instance, your assertion that London is the capital of the UK is a good one, whereas my assertion that my apartment is haunted by ghosts isn't. According to another popular thesis, whether or not an assertion is epistemically good depends on whether or not it satisfies C:

THE C RULE ACCOUNT OF GOOD ASSERTION. One's assertion that p is (epistemically) good if and only if it satisfies THE C RULE OF ASSERTION.

It is easy to see that THE C RULE OF ASSERTION and THE C RULE ACCOUNT OF GOOD ASSERTION together entail the following:

*Institute of Philosophy, KU Leuven, Kardinaal Mercierplein 2, BE–3000 Leuven, Email: christoph.kelp@hiw.kuleuven.be

THE C ACCOUNT OF GOOD ASSERTION. One's assertion that p is (epistemically) good if and only if p has C.

This paper will take a closer look at accounts of the normativity of assertion that accept THE C RULE OF ASSERTION as well as THE C RULE ACCOUNT OF GOOD ASSERTION and derive THE C ACCOUNT OF GOOD ASSERTION from these two theses. Note that any such account is *rule first* in the sense that the conditions for epistemically good assertion are derived from the epistemic rule of assertion. Accordingly, in what follows, I will refer to any such account as a 'rule first account of assertion' (or 'RFAA' for short).

While many researchers are happy to accept RFAA, there is a live debate over the crucial property C. According to perhaps the most popular view, C is identified with knowledge. Champions of a knowledge version of RFAA (henceforth also 'RFAA-K') accept:

THE KNOWLEDGE RULE OF ASSERTION (K RULE). One must (epistemically): assert p only if one knows p.[1]

THE K RULE ACCOUNT OF GOOD ASSERTION. One's assertion that p is (epistemically) good if and only if it satisfies K RULE.

From K RULE and THE K RULE ACCOUNT OF GOOD ASSERTION, champions of RFAA-K derive the following thesis on good assertion:

THE K ACCOUNT OF GOOD ASSERTION. One's assertion that p is (epistemically) good if and only if one knows that p.

Support for RFAA-K is said to derive from assertions of lottery propositions (e.g. 'Your ticket won't win the lottery') and Moore-paradoxical propositions (e.g. 'It is raining but I don't know that it is'). Intuitively, these assertions are not good assertions. Since we are not in a position to know lottery and Moore-paradoxical propositions, RFAA-K is well positioned to explain these intuitions: such assertions violate THE K ACCOUNT OF GOOD ASSERTION. Further support for RFAA-K derives from a family of linguistic data including the fact that assertions can normally be challenged by asking how the speaker knows and won't remain in good standing

---

[1] Champions of K RULE include Peter Unger [1975], Michael Slote [1979], Keith DeRose [2002], John Turri [2011, 2014] and most famously Timothy Williamson [1996, 2000]. K RULE's main rival, which does not lag far behind in popularity, identifies C with a slightly weaker epistemic property, *viz.* justification (alternatively rational credibility, reasonableness to believe, etc.) [Douven 2006, 2009, Lackey 2007, 2008, Gerken 2012, McKinnon 2013, 2015] or justified belief [Kvanvig 2009, 2010]. Slightly less popular, at least among epistemologists, are the truth rule [Weiner 2005, Whiting 2013] and the belief rule [Bach 2008, Hindriks 2007].

Another issue is whether the epistemic rule of assertion is context-sensitive in the sense that what is required for epistemically permissible assertion varies across contexts. Most champions of context-sensitivity opt for a version of the justification rule of assertion and maintain that the degree of justification needed for epistemically permissible assertion varies with context (Gerken, McKinnon and, arguably, Lackey). A noteworthy exception is [Goldberg 2015] according to whom knowledge is often required for epistemically permissible assertion and yet on occasion epistemically permissible assertion may require less or more than knowledge.

unless the speaker can adequately respond to this kind of challenge, which suggests that assertion is indeed governed by K RULE [e.g. Unger 1975, Williamson 2000, Turri 2014].

The first aim of this paper (§1) is to argue that RFAA-K is right in at least one respect: K RULE is true.[2] In the remainder of the paper I will develop in some detail a novel account of the normativity of assertion and argue that it compares favourably with RFAA-K. More specifically, I will first develop an account of etiological functions and their normative import (§2). §3 argues that assertion has the etiological function of generating knowledge in hearers and develops the proposed account of the normativity of assertion by unpacking the normative import of this thesis. Given the primacy of etiological functions in my account, I will henceforth refer to it as 'the function first account of assertion' or 'FFAA' for short. Finally, in §4, I compare RFAA-K and FFAA and show that there is reason to believe that FFAA is the more attractive alternative.

1 The Knowledge Rule Actually Governs Assertion

Consider any performance type that is governed by a rule to the effect that one must: ϕ only if one meets condition C. What I would like to suggest is roughly the following: When a performance type is governed by such a rule, we can *legitimately criticise* agents for ϕing without satisfying C. Here is a slightly more careful statement of this idea:

LEGITIMATE CRITICISM. For any performance type, ϕ, that is governed by a rule to the effect that one must: ϕ only if one meets C, *if we criticise an agent who ϕs without satisfying C for ϕing without satisfying C, our criticism will be **prima facie legitimate***. At the very least, it will be prima facie legitimate when we also satisfy the following two conditions: (i) we know that the agent ϕs but doesn't satisfy C and (ii) we can offer a reason to believe that the agent does not satisfy C when it isn't

---

[2] [Author 2015a]. It may be worth noting that the argument also targets contextualists about the rule of assertion (see fn.1) who accept that, in certain contexts, epistemically permissible assertion requires less than knowledge. At the same time, it does not affect contextualists who hold that, while epistemically permissible assertion always requires knowledge, in certain contexts, it requires more than knowledge. Incidentally, I take it that knowledge is not only necessary but also sufficient for epistemically permissible assertion (see [Simion 2015] for what I take to be a compelling defence of this claim). While, strictly speaking, then, this paper leaves room for at least a certain kind of contextualism about the rule of assertion, I believe that there is good independent reason to avoid going down that route.

obvious that s/he doesn't.[3,4]

That criticisms are legitimate in such cases is, in turn, evidenced by the fact that, in these cases, our criticisms call for or at least make appropriate certain kinds of response on the part of violators, including apologies, excuses, explanations, attempts at rectifications and so on. Consider, for instance, the performance type of playing one's penultimate card in a game of Uno. This performance type is governed by a rule to the effect that one must: call 'uno' when producing a token of the type. When we know that a player of Uno plays his penultimate card without calling 'uno', we can prima facie legitimately criticise his performance, for instance by saying: 'You didn't call 'uno'!'.[5] This, in turn, is evidenced by the fact that our criticism calls for or makes appropriate certain responses on the part of the violator: for instance, he may say that he is sorry (apology), that he didn't know (excuse), that he didn't call 'uno' because someone threatened to shoot him if he did (explanation) and so on.[6]

Consider next performance types such that there is no rule to the effect that one must: ϕ only if one satisfies C. Here I would like to suggest, roughly, that criticisms of agents who ϕ without satisfying C will be *prima facie illegitimate*. Again, the following is a more careful expression of this idea:

ILLEGITIMATE CRITICISMS. For any performance type, ϕ, that is *not* governed by a rule to the effect that one must: ϕ only if one meets C, *if we criticise an agent who ϕs*

---

[3] Note that, even in cases in which we can criticise agents for ϕing without satisfying C with *ultima facie* legitimacy, not every way of criticising the agent for ϕing need be legitimate. After all, there are ways of criticising the agent for ϕing that do not discriminate between violations of the particular rule and the all-things-considered permissibility of ϕing. 'You shouldn't have ϕed' is a case in point. A more discriminating way of stating our criticism is by saying 'You don't C!' (alternatively: 'You don't ...!' where ... is filled by an expression the semantic value of which is entailed by the addressee's satisfying C). Here it is clear that our criticism targets the particular rule proscribing ϕing without satisfying C.

[4] I don't mean to suggest that violations of rules only ever legitimise the kinds of criticisms mentioned above. On the contrary, other kinds of sanctions may be legitimate or even called for as well (perhaps given that further conditions are in place). Crucially, however, this fact does not affect the prima facie legitimacy of the aforementioned kinds of criticism. Rather, violations of rules always render these kinds of criticisms prima facie legitimate, whilst allowing for further kinds of sanction to exist alongside them.

[5] I take it that, in this kind of case, it will be obvious that our Uno player didn't call 'uno' when he didn't and so there is no need to produce a reason supporting the proposition that he didn't call 'uno'.

[6] There is an important difference between legitimate criticism and blame. In particular, there may be cases in which an agent can be legitimately criticised for violating a certain rule, without thereby being deserving of blame. For instance, our Uno player who doesn't call 'uno' when playing his penultimate card may not be deserving of blame for not calling 'uno'. He won't be if he has an excuse or an overriding reason for not doing so. Even so, when we know he didn't call 'uno' and we criticise him for it, our criticism will be prima facie legitimate. It may also be worth noting that, as a result, certain ways of criticising agents for ϕing may at least not be conversationally appropriate in certain cases because they constitute ways of blaming the agent, if only indirectly. 'You shouldn't have ϕed.' is, again, a case in point. (See [Author 2015a] for an application to the case of assertion of a general normative framework for prima facie rules that can be broken blamelessly, which also provides a unified and principled account of the classical putative counterexamples to K RULE.)

*without satisfying C for ϕing without satisfying C, our criticism will be **prima facie illegitimate***. Moreover, our criticism will be prima facie illegitimate, even when (i) we know that the agent ϕs without meeting C and (ii) we can offer a reason that the agent doesn't satisfy C when it isn't obvious that s/he doesn't.

For instance, there is no rule of Uno to the effect that one must call 'duo' when playing one's antepenultimate card. Now suppose an Uno player played his antepenultimate card without calling 'duo' and we know this. It is hard to deny that any criticism levelled against this agent for playing his antepenultimate card without calling 'duo' will be prima facie illegitimate.

By way of further evidence for this, notice that, unless there is special reason which renders our criticism legitimate after all, a direct rebuttal of the criticism as irrelevant or misplaced will be prima facie legitimate. Unless there is special reason rendering our criticism of the Uno player for playing his antepenultimate card without calling 'duo' legitimate, the Uno player can prima facie legitimately rebut our criticism as misplaced or irrelevant. 'So what?' is a prima facie legitimate response to 'You didn't call 'duo'!', unless there is special reason rendering the criticism legitimate. Or, again, more carefully:

LEGITIMATE REBUTTALS. For any performance type, ϕ, that is *not* governed by a rule to the effect that one must: ϕ only if one meets C, *if we criticise an agent who ϕs without satisfying C for ϕing without satisfying C, the agent **may prima facie legitimately rebut** our criticism as misplaced or irrelevant*. At the very least the agent may do so if there is no special reason rendering our criticism legitimate. What's more, the agent may do so even when (i) we know that the agent ϕs without satisfying C and (ii) we can offer a reason that the agent doesn't satisfy C when it isn't obvious that s/he doesn't.

To mount a case for K RULE, I will argue first that the predictions K RULE makes when combined with LEGITIMATE CRITICISMS are correct and, second, that the assumption that K RULE is false leads to incorrect predictions when combined with ILLEGITIMATE CRITICISMS and LEGITIMATE REBUTTALS.

Suppose a speaker asserts that p but doesn't know that p. In this case, K RULE combined with LEGITIMATE CRITICISMS predicts that we can prima facie legitimately criticise the speaker for his assertion, for instance, by saying: 'You don't know that!'. At least, we may so criticise the speaker when (i) we know that he doesn't know that p and (ii) we can offer a reason that the speaker doesn't know that p when it isn't obvious that he doesn't.

Since in this kind of case it will often not be obvious that the speaker doesn't know that p, we will do well to set up our test cases such that we can offer a reason of the relevant kind. Every case in which a speaker doesn't know that p falls in at least one of the following four categories: the speaker's environment is epistemically inhospitable for p; he doesn't have knowledge-level justification for p; p is false; or he doesn't believe that p. And, conversely, every case that falls in at least one of these four categories is a case in which the speaker doesn't know that p. Given that this is so, one way of offering a reason why the speaker doesn't know

that p is by pointing out that the speaker falls in one of these four categories of case. This can be achieved either by asserting 'You don't know that!' and then, by way of explanation, adding the reason why the agent doesn't know (e.g. 'There are too many indistinguishable fakes in the area.' to indicate that the agent is in an epistemically inhospitable environment); or else by offering a more specific form of criticism, which already wears the relevant reason on its sleeve (e.g. 'That's false!', 'You don't believe that yourself!' or 'You have no reason to believe that!').

When we know that the speaker doesn't know that p, then, K RULE combined with LEGITIMATE CRITICISMS predicts that we can prima facie legitimately criticise him for his assertion by saying 'You don't know that p!' with an addendum of a reason why he falls into one of the four categories of cases or else by producing one of the more specific criticisms just mentioned.

There is excellent reason to believe that K RULE's predictions are correct. To see this, notice that, in all of these kinds of case, criticisms call for or at least make it appropriate for the speaker to apologise, offer an excuse or explanation, attempt to rectify the situation and so on. When it is pointed out that the speaker is in an inhospitable environment, he may offer an excuse: 'Oh, I didn't know that.'; when the speaker is criticised for not having knowledge-level justification, he may try to rectify the situation by weakening his assertion: 'Okay, okay, but there is good reason to believe that p.'; when the criticism concerns the truth-value of the assertion, the speaker may offer an explanation: 'You're right. But that guy threatened to shoot me unless I said this.'; when she is admonished for not believing what she asserted, she may offer an apology: 'I am sorry, I misspoke.'.[7]

Let's now assume that assertion is not governed by K RULE. By ILLEGITI- MATE CRITICISMS, we get the prediction that criticisms levelled against speakers for asserting what they don't know will be prima facie illegitimate, even when their critics know that they don't know what they assert and can offer a reason why they don't know. This prediction is manifestly incorrect. As the discussion above clearly indicates, when we know that a speaker doesn't know what he asserts and can of- fer such a reason, our criticisms of him for asserting what he doesn't know will be prima facie legitimate. If so, our criticisms will not be prima facie illegitimate.

In addition, assuming that assertion is not governed by K RULE in conjunction with LEGITIMATE REBUTTALS leads to the prediction that, unless there is special reason rendering criticisms of speakers for asserting what they do not know legiti- mate after all, speakers can prima facie legitimately rebut the criticisms as irrelevant or misplaced, even when their critics know that they don't know what they assert and can offer a reason why they don't know. Again, this prediction is manifestly in- correct. Even when there is no special reason rendering such criticisms legitimate, speakers cannot prima facie legitimately rebut them as irrelevant or misplaced. 'So what?' is not a prima facie legitimate response to 'You don't know that!' and so on, at least not given that we know that the speaker doesn't know what he asserts and

---

[7] I do not mean to suggest that the various kinds of case invariably call for the kind of response described above. Different situations may make different responses appropriate. What matters is that some response in the range is always called for/made appropriate.

have offered a reason why he doesn't know.[8]

Since K RULE in conjunction with LEGITIMATE CRITICISMS makes correct predictions and the assumption that K RULE is false in conjunction with ILLEGITIMATE CRITICISMS or LEGITIMATE REBUTTALS leads to incorrect predictions, there is reason to believe that assertion is governed by K RULE.

2 Etiological Functions and their Normative Import

2.1 Etiological Functions

I would now like to investigate the relation between etiological functions (henceforth also 'e-functions' for short) and certain kinds of norms. To begin with, let's look at the following plausible account of e-functions:

> E-FUNCTION. A token of type T has the e-function of producing effect E in system S iff
>   (EF1) Past tokens of T produced E in S's ancestors
>   (EF2) Producing E benefitted S's ancestors
>   (EF3) Producing E's having benefitted S's ancestors contributes to the explanation of why T exists in S.[9]

Consider your heart, which plausibly has the e-function of pumping blood. Notice that one effect that past tokens of hearts had in your ancestors was to pump blood (EF1). Moreover, this effect was beneficial to your ancestors (EF2). In fact it was of vital importance. If their hearts had not pumped blood, they would not have survived very long. Finally, that past hearts pumped blood in your ancestors contributes to the explanation of why hearts exist in present-day humans (EF3). Past hearts pumping blood contributed to the proliferation of the genes responsible for them and in this way to the explanation of the existence of the heart in present day humans. By E-FUNCTION, your heart has the e-function of pumping blood.

In contrast, while your heart also produces a certain kind of sound, producing this kind of sound is not an e-function of your heart. One reason for this is that

---

[8] It may be worth considering the predictions of K RULE's most prominent competitor in the literature, the justification rule. In conjunction with ILLEGITIMATE CRITICISMS, the justification rule predicts that when, in the cases just described, we criticise the agent accordingly, our criticism will be prima facie illegitimate. For instance, 'That's false!' turns out to be a prima facie illegitimate criticism, even when we know that what the speaker asserted is false. What's more, unless there is special reason rendering the criticism legitimate, in conjunction with LEGITIMATE REBUTTALS, the justification rule predicts that the criticism can be prima facie legitimately rebutted as irrelevant or misplaced. 'So what?' can be a prima facie legitimate rejoinder to 'That's false!'. It is easy to see that these predictions are manifestly incorrect.

[9] Etiological accounts of functions are the most widely accepted accounts in contemporary literature. Prominent champions include Ruth Millikan [e.g. 1984], Karen Neander [e.g. 1991b, 1991a], Paul Griffith [1993], Peter Godfrey-Smith [1993, 1994] and David Buller [1998]. That said, there are a number of noteworthy dissenters, including John Bigelow and Richard Pargetter [1987], Philip Kitcher [1994] and Denis Walsh and André Ariew [1996]. For applications to epistemology see e.g. [Graham 2012, 2014b,a]. In fact, E-FUNCTION is close to Peter Graham's [2014b] account of e-functions. It differs from Graham in explicitly opting for a weak account of e-functions along the lines of [Buller 1998].

producing the relevant kind of sound had no discernible benefit for your ancestors. (EF2) is not satisfied. Even if we assume (EF2) to be satisfied, say because the kind of sound hearts produced was pleasant to your ancestors, production of this benefit does not contribute to the explanation of why hearts exist in humans. So (EF3) is not satisfied in any case. By E-FUNCTION, your heart does not have the e-function of producing this kind of sound.[10]

While your heart is a paradigm case of an item with an e-function, I'd also like to look at a slightly more complex case (or, perhaps, a slightly different case with a bit more explicit structure). Consider a system that is constituted by the following types of subsystem: a type of producer, a type of product, a type of consumer, and a type of return such that the producer produces token products, which may be consumed by consumers in exchange for a return. In what follows, I will refer to any such system as a *Simple Economic System (SES)*. By way of example of a SES, consider the practice of producing an consuming cups of espresso at your local coffee shop (SES-Espresso). This practice is a system involving baristas (producers) who produce cups of espresso (product). Furthermore the system features costumers at the coffee shop (consumers), who may purchase (return) and then consume tokens cups of espresso.

Tokens of the product in an SES may (and often will) have the e-function of producing a certain effect, E, in consumers. One might wonder how this could be given that according to E-FUNCTION, the effect must occur in the system, i.e. the SES. Recall, however, that the consumer is a subsystem of a given SES that partly constitutes it. Since, plausibly, an effect in a system may be produced by producing it in one of its constituent subsystems, the idea that tokens products may (and often will) have the e-function of producing E in consumers should not raise too many eyebrows.

By E-FUNCTION, token products will have the e-function of producing E in consumers if and only if past tokens of the product produced E in past consumers (EF1), E was beneficial for an ancestor (i.e. here: a past self) of the SES (EF2) and the fact that E was beneficial for an ancestor of the SES contributes to explaining why the SES continues to exist today (EF3).[11]

---

[10] I borrow this example from [Graham 2012: 9].

[11] This is a somewhat simplified formulation of EF3, which should really read as follows: the fact that E was beneficial for an ancestor of the SES contributes to explaining *why the type of product continues to exist* in the SES today. Note, however, that when the SES features only a single type of product, the SES will cease to exist when and only when the product ceases to exist in the SES. For instance, assuming that cups of espresso are the only type product produced by baristas in SES-Espresso, SES-Espresso will cease to exist when and only when cups of espresso cease to exist at your local coffee shop. Of course, this is not the case in SESs that feature more than one type of product. Consider, for instance, the practice of producing and consuming coffee products at your local coffee shop, including espressos, cappuccinos, lattes and so on. Here the relevant SES may continue to exist even if cups of espresso don't continue to exist in it. Accordingly, EF3 cannot be simplified in the above way. Since, for the purposes of this paper, I will be focusing only on SESs that feature a single type of product, the simplification is unproblematic. Accordingly, the fact that I use it frequently in what follows should not be cause for concern either.

I take it that it is pretty clear how EF1 can be satisfied in an SES. For instance, in SES-Espresso, past cups of espresso may very well have generated pleasant gustatory experiences in customers. Assuming that they did, EF1 is satisfied here. The cases of EF2 and EF3 might benefit from some further commentary. EF2 may be satisfied in a given SES in virtue of the fact that E was beneficial for past consumers. In that case, the benefit for the SES resides in the benefit for consumers.[12] In SES-Espresso, for instance, it is plausible that the production of pleasant gustatory experiences in customers is beneficial for SES-Espresso in virtue of being beneficial for customers. The benefit of pleasant gustatory experiences in customers for SES-Espresso resides in the benefit of such experiences for customers here. Since I take it to be clear enough that pleasant gustatory experiences do constitute a benefit for customers, EF2 will be satisfied. Finally, EF3 will often be satisfied in a given SES because the beneficial effect produced in past consumers of past tokens of the product motivated them to continue purchasing tokens, which, in turn, motivated producers to continue producing tokens of the type, thus contributing to the explanation of why the SES continues to exist. For instance, in SES-Espresso, the fact that token cups of espresso generated pleasant gustatory experiences in costumers motivated costumers to continue purchasing cups of espresso, which, in turn motivated baristas to continue producing them, thus contributing to explaining why SES-Espresso continues to exist today.

2.2 Two Kinds of Norm

With this account of e-functions in play, I'd now like to introduce a distinction between two kinds of norm, which I take from Conor McHugh. McHugh characterises these two kinds of norm as follows:

> PRESCRIPTIVE NORMS. "[Prescriptive norms] are to do with what one ought, may or ought not do: they require, permit or forbid certain pieces of conduct on the part of agents, and are apt to guide that conduct." [McHugh 2012: 9]

> EVALUATIVE NORMS. "[Evaluative norms] are norms primarily to do with what is good or bad, valuable or disvaluable." [McHugh 2012: 10]

Examples of prescriptive norms include many moral norms, such as the norm forbidding stealing, and traffic norms, such as the norm requiring drivers to stop for at least 3 seconds at a stop sign. It is clear that these norms, respectively, forbid and require certain forms of conduct and are apt to guide an agent's behaviour. It is easy to see that what I have so far called 'rules' are prescriptive norms in this sense. Accordingly, for the sake of terminological consistency and ease of exposition, in what follows I will use 'rule' and 'prescriptive norm' interchangeably.

---

[12] I do not mean to suggest that the benefit to the SES must always reside in the benefit for consumers. On the contrary, it seems plausible that, on occasion, there may be a benefit for the SES without there being a benefit for consumers. Conversely, it also seems plausible that a benefit for consumers may not be a benefit for the SES. All I am saying is that the benefit for the system may reside in the benefit for the customer.

In contrast, evaluative norms are not prescriptive in this way. They don't tell agents what to do. Rather, they specify conditions of attributive goodness, i.e. goodness as it pertains to a certain kind.[13] There is such a thing as a good knife or a good sprinter. An evaluative norm specifies the conditions under which a knife is a good knife or a sprinter is a good sprinter. 'A good knife is a sharp knife' would be an example of an evaluative norm, 'a good sprinter runs fast' another. (Note that THE K ACCOUNT OF GOOD ASSERTION also qualifies as an evaluative norm in this sense.)[14]

Finally, as McHugh also points out, while there may be relations between prescriptive and evaluative norms, these are by no means straightforward. In what follows, I want to take a closer look at one way in which e-functions, evaluative norms and prescriptive norms are related. More specifically, I will argue (i) that e-functions give rise to evaluative norms and (ii) that, under certain conditions, e-functions serve to rationalise prescriptive norms as well.

2.3 Etiological Functions and Evaluative Norms

In what follows I will argue that once a token of a type has an e-function, there are facts about what it is for a token to be a good token of its type. In this way, e-functions have normative import. To get there, I will first have to introduce some further conceptual machinery.

The first concept we need is the concept of function fulfilment. For any token that has an e-function, there is such a thing as function fulfilment. Roughly, a token with an e-function fulfils its function if and only if it produces the relevant beneficial effect. For instance, given that the e-function of a cup of espresso in SES-Espresso is to produce pleasant gustatory experiences in customers, a token cup of espresso fulfils its e-function if and only if it generates pleasant gustatory experiences in some costumer.

Second is the concept of normal functioning. For any token that has an e-function, there is such a thing as normal functioning. This is the way in which past tokens functioned when the type acquired its e-function. That is to say, it is the way in which past tokens of the type functioned when producing the benefit that made the contribution to the explanation of why the type now exists in the system. In the case of SES-Espresso, it is the way past cups of espresso functioned when they generated pleasant gustatory experiences in past costumers. Let's say that this includes the stimulation of certain taste buds in the mouths of consumers, etc.

---

[13] The notion of attributive goodness is due to Peter Geach [1956]. Characteristic of this attributive sense of 'good' is that the reference of 'good K' does not divide into 'is good' and 'is a K'. Note that Geach also claims all uses of 'good' are attributive. In particular, there is no such thing as good simpliciter, independently of a kind. I want to stay neutral on this latter issue. What matters is that there is such a thing as attributive goodness, which seems safe enough.

[14] It may be worth noting that evaluative norms can also be expressed in terms of ought. The evaluative norm for knives and can be put as follows: 'Knives ought to be sharp' and 'Sprinters ought to be fast'. Note, however, that the oughts here are *ought-to-bes*. Unlike the *ought-to-dos* at issue in prescriptive norms, they do not require, permit or forbid pieces of conduct.

The third and last concept is the concept of normal conditions. For any token that has an e-function there is such a thing as normal conditions. These are the conditions in which past tokens were situated when the e-function was acquired. That it to say, it is the conditions in which past tokens of the type produced the benefit which made the contribution to the explanation of why the type now exists in the system. Let us say that, in the case of SES-Espresso, this includes that token cups of espresso were properly consumed (drunk rather than snorted), that consumers had properly functioning tastebuds, that the air was sufficiently clean, etc.

With these points in play, here is the key thesis concerning the normative import of e-functions. E-functions give rise to evaluative norms: a token of type T with the e-function of producing E may be a good token of T. Crucially, this may be the case even if the token does not fulfil its e-function. A token of cup of espresso that is forgotten on the counter may be a good cup of espresso (at least for a while), even though it is not fulfil its function. Why is that? I would like to suggest the following answer:

THE EVALUATIVE NORM OF TOKENS WITH AN E-FUNCTION. A token of type T with the e-function of producing E is a good T if and only if it has the disposition to fulfil its e-function by functioning normally when in normal conditions.

THE EVALUATIVE NORM OF TOKENS WITH AN E-FUNCTION accommodates the datum that, in the above case, the cup of espresso is a good cup of espresso. After all, even though it has been forgotten on the counter, it retains the disposition to fulfil its e-function by functioning normally when in normal conditions. That is to say, it retains the disposition to produce pleasant gustatory experiences (function fulfilment) by stimulating tastebuds, etc. (normal functioning) when properly consumed, tastebuds are functioning properly, the air is sufficiently clear, etc. (normal conditions).

In contrast, suppose that a token cup of espresso is spoiled by a trickster who has added salt to the espresso. In this unfortunate situation, the cup of espresso no longer has the disposition to fulfil its e-function by functioning normally when in normal conditions. THE EVALUATIVE NORM OF TOKENS WITH AN E-FUNCTION predicts, again correctly, that the cup of espresso is not a good cup.

2.4 E-Functions and Prescriptive Norms

Thus far I have developed a general account of e-functions and have argued that these e-functions have normative import. In particular, they give rise to evaluative norms for tokens with an e-function. But now recall that, besides evaluative norms, there are also prescriptive norms. The question thus arises how prescriptive norms fit into the e-functionalist framework. To answer it, I will now look at a certain kind of situation in which it makes sense to govern the production of tokens with an e-function by prescriptive norms.

To get started, let's suppose that tokens of a certain type have a certain e-function. Now suppose that, in addition, the following conditions obtain:

RELIABILITY. It matters that tokens of the type fulfil their e-function reliably.

VARIATION. There are a number of ways of producing tokens of the type which differ in the degree of reliability with which the tokens produced fulfil their e-function.

HUMANITY. Tokens of the type are produced by human agents who are both free and criticism-averse.

In situations with these features, it makes sense to regulate the production of tokens with an e-functional type by *rules* that prescribe the use of ways of producing tokens of the relevant types that reliably lead to the fulfilment of the relevant e-function. Since tokens are produced by free but criticism-averse human agents and since violating the rules will make producers liable to criticism, regulating production by such rules will lead producers to favour using prescribed ways of producing tokens, at least given that all else is equal. Given that prescribed ways reliably lead to the fulfilment of the relevant e-function, this in turn, will increase (or sustain) the ratio of tokens that fulfil their relevant e-functions to those that don't. In this way, regulating production of tokens with the e-function by a rule will contribute to ensuring that tokens reliably fulfil the relevant e-function. Since this is something of value, it makes sense to regulate production of tokens with an e-function in situations with the above features by rules.

By way of illustration, consider once again our toy case of SES-Espresso. Let us suppose that it matters that cups of espresso produce pleasant gustatory experiences reliably here (RELIABILITY), because, otherwise, baristas will no longer make a profit off them and will stop producing them. There are many ways of producing cups of espresso, which differ in the degree of reliability with which cups of espresso produced by them produce pleasant gustatory experiences in costumers (VARIATION). One may produce them by filling the portafilter with unground coffee beans, the espresso machine with juice and the cup with exactly three millilitres of liquid and so on (W1). Alternatively, one may produce them by filling the portafilter with ground coffee, the espresso machine with water and the cup with approximately 25 millilitres of liquid (W2) and so on. It is not hard to see that producing cups of espresso via W1 will yield cups of espresso that produce pleasant gustatory experiences in customers with a very low degree of reliability, whereas W2 achieves a higher degree of reliability. Cups of espresso at the local coffee shop are produced by baristas, that is, by human employees who are free in the sense that they could produce cups of espresso via a variety of ways, including W1 and W2, but also criticism-averse and so will, all else equal, seek to avoid ways of producing cups of espresso that make them liable to criticism (HUMANITY).

In this case, it makes sense to regulate production of cups of espresso at the local coffee shop by a set of rules, including a rule that requires employees to produce espressos via W2. Once the rule is in place, any way of producing espressos other than W2 will make employees liable to criticism. Since employees are human agents and so criticism-averse, the rule will lead employees to favour producing cups of espresso via W2 over other ways of producing espressos, including W1. This, in turn, will increase (or sustain) the ratio of token cups of espresso that are

produced by W2, a reliable way of producing cups of espresso that produce pleasant gustatory experiences in customers, and so contribute to ensuring that cups of espresso reliably produce pleasant gustatory experiences in customers at the coffee shop. Since it matters that cups of espresso produce pleasant gustatory experiences in customers reliably, it makes sense to regulate production of them by the rule in question.

Finally, suppose that, in some case in which it makes sense to govern production of tokens with an e-function by a rule because it matters that token products fulfil their e-function reliably, we also have independent reason to believe that production of tokens of a certain type of product is actually governed by a particular rule. For instance, suppose that, in SES-Espresso, there actually is a rule that requires producers to produce cups of espresso via W2. If it can be shown that governing production of the relevant tokens by the particular rule does actually contribute to ensuring that they reliably fulfil their e-function, then we will have a rationalisation of this rule. For instance, if it can be shown that regulating production of cups of espresso in SES-Espresso by a rule that requires producers to produce them via W2 does contribute to ensuring that cups of espresso reliably produce pleasant gustatory experiences in customers, then we will have a rationalisation of this rule. It makes sense not only to govern production of cups of espresso by some rule, it makes sense to govern it by *this rule in particular*.[15]

2.5 Typing Norms

I want to suggest that e-functions can be typed in accordance with the type of benefit that they produced in the ancestors of its possessor and that contributed to explaining why the system continued to exist. For instance, in the case of your heart, the relevant benefit is survival, i.e. a biological benefit. Accordingly, the e-function of your heart is a biological e-function. In contrast, in the case of cups of espresso in SES-Espresso, the relevant benefit is a pleasant gustatory experience, i.e. a culinary benefit.

Similarly, evaluative norms and rules that e-functions of a certain type give rise to can also be typed in accordance with the type of benefit they produced in the ancestors of its possessor and that contributed to explaining why the system continued to exist. For instance, in the case of your heart, the evaluative norm that derives from your heart's biological e-function of contributing to survival is a biological norm. It states what it is for your heart to be a biologically good heart. In the case of cups of espresso in SES-Espresso, the evaluative norm that derives from cups of espresso's producing pleasant gustatory experiences is a culinary norm. It states

---

[15] This is of course not to say that the rule we have thus rationalised is the best rule we could operate. There may well be room for improvement. For instance, there may be alternative rules that will do an even better job at ensuring that token products fulfil their e-function reliably. In this case, we may have reason for changing rules. While such reason may in certain cases defeat the rationalisation, it needn't always do so. It won't when we (non-culpably) fail not know of the existence of any better rules, when the costs associated with implementing a change of rules do not match the benefits that arise from increase in reliability and so on. I take these considerations to show that the kind rationalisation of an existing rule sketched above really is a prima facie rationalisation.

what it is for a cup of espresso to be a culinarily good cup of espresso. Similarly, any rule governing production of cups of espresso in SES-Espresso that serves to ensure that cups of espresso fulfil their culinary function reliably will be a culinary rule.

A given token of a type may simultaneously possess several types of e-function. Token cups of espresso in SES-Espresso may have the culinary e-function of generating pleasant gustatory experiences in customers. However, they also have the economical e-function of making a profit for baristas.

Different types of e-function give rise to different types of function fulfilment, norms and rules. For instance, a cup of espresso in SES-Espresso will fulfil its culinary e-function if and only if it generates pleasant gustatory experiences in some customer. It will be a culinarily good cup of espresso if and only if it has the disposition to do so. In contrast, it will fulfil its economic function if and only if it makes a profit for the local coffee shop. It will be an economically good cup of espresso if and only if it has the disposition to make a profit.

Finally, where a token has various e-functions, it can fulfil its function with respect to one of them without doing so with respect to another and satisfy one norm/rule generated by one without satisfying the norm/rule generated by the other. If a cup of espresso in SES-Espresso is spilled right after purchase, it fulfilled its economic but not its culinary e-function. On the other hand, a cup of espresso that has been spoiled by salt and, as a result, is not a culinarily good cup of espresso may still have the disposition to make a profit for the local coffee shop, thus being a good cup of espresso economically. And it may of course still make a profit for baristas, thus fulfiling its economic e-function.

3 The Normativity of Assertion

3.1 The E-Function of Assertion

A number of contributors to the literature [e.g. Millikan 1984, Graham 2010] have argued that assertion has the epistemic e-function of generating true belief in hearers. In what follows, I will first reconstruct the argument in a way that lines up with some of the earlier discussion about SESs.

Linguistic practices, including practices of producing and consuming linguistic devices such as assertions, can be viewed as SESs. Producers are speakers, linguistic devices such as assertions are products, and consumers are hearers. This much is easy.

However, SESs also feature a type of return. What is the return in the case of linguistic practices in general and in the case of assertion in particular? To answer these questions, consider first the following passages from Millikan:

> Language devices will produce effects that interest speakers often enough to encourage continued replication only if hearers replicate hoped-for cooperative responses often enough.

> [Millikan 2004: 25]

14

What comes to light here is that linguistic practices do feature a return, *viz.* a *co-operative response* on the part of the hearer. Moreover, taking another leaf from Millikan and Graham, I want to suggest that, in case of the practice of producing and consuming assertions, the cooperative response is belief on the part of the hearer.[16]

Given that the practice of producing and consuming assertions can be viewed as an SES (henceforth also 'SES-Assertion'), we may now investigate whether the product, i.e. assertion, does indeed have the epistemic e-function of generating true belief in hearers. It will do so if and only if (EF1) past assertions generated true beliefs in hearers, (EF2) this constituted a benefit for SES-Assertion and (EF3) that this constituted a benefit for hearers contributes to the explanation of why SES-Assertion continues to exist.

EF1 is a highly plausible empirical claim. Concerning EF2, note first that, according to a widely accepted thesis, our epistemic aim is to amass a large body of beliefs with a favourable truth/falsity ratio.[17] The generation of true belief in hearers, then, means that hearers make progress in the direction of attaining their epistemic aim. This, I take it, is a benefit for hearers. Note also that it is an epistemic such benefit. Finally, just as in SES-Espresso earlier, it is plausible that the benefit for the system resides in the benefit for the consumer. If so, the generation of true belief in hearers not only constitute a benefit for hearers but also for SES-Assertion as a whole. EF2 is satisfied also.

This leaves us with is EF3. In a nutshell, here is the Millikan/Graham argument: speakers continue to produce a linguistic device only if this often enough produces the cooperative response on the part of hearers. Since, in the case of assertion, the desired hearer response is belief, hearers will respond cooperatively to the production of a linguistic device only if they benefit from responding in this way. In the case of assertion, gaining a true belief is the main benefit for hearers. This is evidenced by the fact that if hearers had not gained true belief sufficiently often by responding to assertions with belief, they would before long have stopped responding to assertion with belief. Since if they had done so, they would no longer produce the cooperative response, before long speakers would no longer have been motivated to make assertions and the practice of producing and consuming assertions would have been discontinued. In this way, the fact that past assertions generated true belief in hearers does contribute to explaining why SES-Assertion continues to exist today. EF3 is satisfied as well.

Since all of EF1 – EF3 are satisfied, it follows that assertion has generating true belief in hearers as an e-function. Moreover, since the benefit that contributes to

---

[16] [Millikan 2004: 26, Graham 2010: §4]. It may be worth noting that Millikan states her argument in terms of "indicative sentences". Even so, it is easy to see that the argument will work just as well for assertion, not in the least because we typically make assertions by means of indicative sentences.

[17] I take it that the thought here is that amassing a large body of beliefs with a favourable truth/falsity ratio is the constitutive aim of the epistemic domain. However, for present purposes, not much hinges on whether the aim is thus constitutive so long as it is indeed our epistemic aim in some sense.

explaining the continued existence of SES-Assertion is an epistemic benefit, generating true belief in hearers is an epistemic e-function of assertion.

I find the above argument by and large compelling. My only disagreement with Millikan and Graham is that I think that the e-function of assertion is stronger than they make out. More specifically, I think it consists in generating *knowledge* in hearers, not just *true belief*.

Holding on to the first part of the argument, according to which the practice of producing and consuming assertions can be viewed as an SES in which assertion is the product and belief on the part of the hearer the return, what needs to be shown is that past assertions not only generated true belief in hearers but also knowledge (EF1), that this constituted an epistemic benefit for SES-Assertion (EF2), and that the fact that it did so hearers contributes to the explanation of why SES-Assertion continues to exist.

EF1 is a highly plausible empirical claim. Since knowledge entails true belief, it would seem that EF2 remains plausible for the same reasons that motivated the corresponding claim about true belief.

Before turning to EF3, I briefly want to address the question whether this will do by way of support for EF2. After all, if knowledge and not just true belief is to be the e-function of assertion, shouldn't the benefit at issue in EF2 be one that attaches specifically to knowledge? Fortunately, it is quite plausible that knowledge does come with specific benefits, even epistemic ones. After all, as most epistemologists agree, knowledge is epistemically more valuable than true belief.[18] Given that this is so, knowledge comes with epistemic benefits over and above the ones that attach to true belief. EF2 is plausible even if we require that the benefits at issue in EF2 attach to knowledge in particular.

Again, EF3 is the hard one to justify. It may seem as though, in the present case, it is particularly tricky because it is far from clear that SES-Assertion would cease

---

[18] Offering a detailed account of the additional value of knowledge is of course a philosophically substantive task, one that I cannot hope to undertake within the confines of this paper. That said, one account that I am attracted to employs the thesis that knowledge, but not true belief, is the goal of the epistemic practice of inquiry. Since inquiry is an epistemic practice, then, knowledge is an epistemic goal over and above true belief. Since it's better in general to attain the goal of a practice, i.e. to succeed in it, than to fall short of success, it follows that knowledge is epistemically more valuable than true belief.

Of course, that knowledge is the goal of inquiry is another thesis of philosophical substance, one that I will be no more able to defend in adequate detail within this paper than the thesis that knowledge has value over and above true belief. That said, [Author 2014a] takes on this task and provides two arguments in support of the idea that knowledge is the goal of inquiry. Here is a very brief sketch of one of these arguments: The argument's crucial thesis is that attaining success in a practice will automatically release one from any commitments towards attaining that success in that practice. However, one may attain true belief and not be released from all commitments towards attaining success in inquiry. For instance, if one has been contracted to settle the question whether p, a gettiered true belief that p will not automatically release one from one's contractual commitment to settle the question whether p. When the gettiered true belief has been arrived at via inference from a false premise q and it has become known that q is false, one is still bound by the contract. By way of evidence, note that one's employer can legitimately insist that one has not fulfilled the contract and take steps to get one back to work.

to exist altogether if assertions never produced knowledge. Suppose, for instance, that were assertion not to generate knowledge in hearers, this would be because it is common knowledge that, whenever a speaker produces an assertion, the content passes through a randomiser, which outputs the same content with a probability 0.9 and outputs a false content with a probability of 0.1. In that case, assertion never generates knowledge in hearers. Or, to be more precise, the assertion that p never generates knowledge that p in hearers. At the same time, there is little reason to think that SES-Assertion would cease to exist. Assertion is still too useful a tool to abandon it entirely in this situation.

At the same time, EF3 only requires that the benefit produced by past tokens of a type *contributes to the explanation* of why the system continues to exist. Can it be the case that such a benefit contributes to the explanation of why the system continues to exist even though it is not the case that had tokens of the type not produced the benefit, the system would not have existed? As I am about to argue the answer to this question is 'yes'. In order to achieve this, I will argue that satisfaction of EF3 for products in an SES does not generally require the truth of a counterfactual of the form *had past tokens of the product not produced the benefit, the system would not exist today*. I will then offer an alternative account of what it takes to satisfy EF3 and will argue that in SES-Assertions, knowledge does satisfy EF3.

First, let's return to SES-Espresso. Let's continue to assume that past cups of espresso produced pleasant gustatory experiences in customers and that this constituted a benefit for SES-Espresso. If so, EF1 and EF2 are satisfied. But now suppose that even if past cups of espresso had not produced pleasant gustatory experiences in customers, they would still have produced gustatory experience of mediocre quality in them. Suppose, furthermore, that, whilst constituting a lesser good for customers, experiences of this quality would nonetheless have been good enough to make consumers willing to pay for the provision of cups of espresso. This would have been good for the baristas and would have motivated them to continue producing cups of espresso with the result that SES-Espresso would have continued to exist anyway.

In this case, do cups of espresso in SES-Espresso have the e-function of generating pleasant gustatory experiences in customers? Since we have already seen that EF1 and EF2 are satisfied, this boils down to the question of whether EF3 is satisfied also. That is to say, it boils down to the question of whether the fact that past cups of espresso produced pleasant gustatory experiences in customers contributes to the explanation of why SES-Espresso continues to exist today. Crucially, if in order to do this it must be the case that had past cups of espresso not produced pleasant gustatory experiences in customers, SES-Espresso wouldn't have continued to exist today, the answer to this question is 'no'. After all, in the case under consideration, this counterfactual is false.

Should we accept, then, that cups of espresso in SES-Espresso do not have any e-function at all in this case? That also seems implausible. After all, it looks as though there is some benefit that they produced in customers that contributes to explaining why SES-Espresso continues to exist today. So perhaps the thing to say is that, while cups of espresso do have an e-function in this case, it's not the e-

function of producing pleasant gustatory experiences. Rather, it is some different e-function. Should we perhaps say that, in the present case, the e-function of cups of espresso consists not in producing pleasant gustatory experiences, but in producing at least mediocre ones? Before leaping to an answer, notice that it might be that had cups of espresso not even done that, SES-Espresso would still have continued to exist. This might be because in that case they would have produced a different kind of beneficial effect, say an extreme caffeine rush or the pleasant feeling generated by drinking what one takes to be a cool drink in a cool place. But we surely wouldn't want to say that the e-function of cups of espresso really consists in its caffeine rush or the feeling of coolness. If this isn't immediately obvious note that it may very well be that, at the actual world at which cups of espresso produce pleasant gustatory experiences, they produce none of the beneficial effects they would have produced had they not generated gustatory experiences of at least mediocre quality.

What transpires is the following. Even if it is *not* the case that had cups of espresso not produced pleasant gustatory experiences in the past, SES-Espresso would not have continued to exist, it remains plausible that cups of espresso have some e-function here. After all, it is plausible that they produced some benefit in the system that contributes to explaining why SES-Espresso continues to exist today. The problem is that once we give up on the idea that the relevant benefit produced are pleasant gustatory experiences in customers, on the grounds it doesn't satisfy the relevant counterfactual, it is not clear that we will have any satisfactory alternative candidate left. After all, no alternative candidate that is even remotely plausible may satisfy the counterfactual either. In view of this, it seems wiser to abandon the idea that in order to satisfy EF3, a benefit actually produced by a product in an SES must satisfy the counterfactual that had past tokens of the product not produced it, the SES would not have continued to exist today.

Let's move on to the question of what it takes for a product in an SES to satisfy EF3, i.e. what it takes for a certain benefit of past tokens of a type to contribute to the explanation of why the SES continues to exist. I would like to suggest the following answer: EF3 is satisfied if the benefit was actually produced and actually was part of the instantiation of the causal structure in virtue of which the explanation of why the SES continues to exist comes out true. If so, it is easy to see that the benefit of producing pleasant gustatory experiences in customers does contribute to the explanation of why SES-Espresso continues to exist today: the benefit was actually produced by past cups of espresso and is part of the instantiation of the causal structure in virtue of which the explanation comes out true. That's why past cups of espresso contribute to the explanation of why SES-Espresso continues to exist today. On this account, then, present cups of espresso have generating pleasant gustatory experiences as an e-function in SES-Espresso.

Now, I want to suggest that, just as in SES-Espresso cups of espresso have the e-function of producing pleasant gustatory experiences in customers even if SES-Espresso would have existed had they not produced this benefit, so in SES-Assertion assertions can have the e-function of generating knowledge in hearers even if SES-Assertion would have existed today had they not generated knowledge

in hearers. What matters to whether EF3 is satisfied is whether the relevant benefit was actually produced by the product and that it is actually part of the instantiation of the causal structure in virtue of which the explanation comes out true. And this is the case in both SES-Espresso and SES-Assertion. Just as past cups of espresso actually produced pleasant gustatory experiences in customers, so past assertions actually generated knowledge in hearers. Moreover, just as the benefit of pleasant gustatory experiences is actually part of the instantiation of the causal structure in virtue of which the explanation of why SES-Espresso continues to exist today comes out true, so the benefit of generating knowledge in hearers is actually part of the instantiation of the causal structure in virtue of which the explanation of why SES-Assertion continues to exist today comes out true. Since this suffices for cups of espresso to satisfy EF3 in SES-Espresso, the same goes for assertion in the case of SES-Assertion. If so, assertion does have generating knowledge in hearers as an epistemic e-function.[19]

3.2 The Evaluative Norm of Assertion

Given that assertion has the epistemic e-function of generating the knowledge in hearers, we get the corresponding accounts of what counts as function fulfilment, normal functioning and normal conditions.

In the case of assertions' e-function of generating knowledge in hearers function fulfilment will of course consist in generating knowledge in some hearer.

Normal functioning and normal conditions are defined as expected in terms of, respectively, the way assertion functioned back when it acquired its e-function and the conditions that obtained back then.

Most importantly for present purposes, however, we also get the following evaluative norm:

THE EVALUATIVE NORM OF ASSERTION. One's assertion that p is (epistemically) good if and only if it has the disposition to generate knowledge that p in one's hearer(s) (function fulfilment) by functioning normally when in normal conditions.

3.3 The Prescriptive Norm of Assertion

Recall that we looked at tokens with an e-function such that the following conditions are satisfied:

---

[19] Are there further epistemic e-functions of assertion, besides this one? I am not sure. That said, I will not attempt to answer this question in any detail here. Rather, I simply assume that generating knowledge in hearers is the only epistemic e-function of assertion. While I agree that this is a substantive assumption, it is at least not evident that it will be of much consequence for the purposes of this paper. That is to say, it is not clear that, even if there turned out to be further epistemic e-functions besides this one, the paper's central argument will not go through. For that reason, I will leave it to those who want to criticise FFAA on the grounds that it countenances only one epistemic e-function of assertion to establish that assertion has further epistemic e-functions and to show that this compromises FFAA.

RELIABILITY. It matters that tokens of the type fulfil their e-function reliably.

VARIATION. There are a number of ways of producing tokens of the type which differ in the degree of reliability with which the tokens produced fulfil their e-function.

HUMANITY. Tokens of the type are produced by human agents who are both free and criticism-averse.

I argued that in this case it makes sense to govern production of token with an e-function by a rule. For instance, it makes sense to govern the production of cups of espresso in SES-Espresso by a rule that prescribes a certain way of producing cups of espresso. Moreover, given that production of tokens with an e-function is actually governed by a rule, we can provide a rationalisation of the rule by showing that the rule does contribute to ensuring that token products fulfil their relevant e-function reliably. Given that production of cups of espresso in SES-Espresso is governed by a specific rule, we can rationalise this rule if we can show that operating this rule does contribute to ensuring that cups of espresso produce pleasant gustatory experiences in customers reliably.

What I will argue now is that assertion in SES-Assertion is another case in point. Here is why. First, in the case of assertion, it matters that token assertions fulfil their e-function of generating knowledge in hearers reliably (RELIABILITY). If they didn't, say because more often than not assertion produced beliefs in hearers that are either unjustified, false or gettiered, it is hard to see how assertion could continue to be a source of knowledge at all. Even if this doesn't quite mean that assertion will cease to exist altogether (see §3.1), at the very least SES-Assertion will undergo a significant transformation: its e-function will change. Moreover, there are many ways of producing assertions which differ in the degree of reliability with which assertions produced by them generate knowledge in hearers (VARIATION). One may assert that p based on a mere guess that p (W1), an educated guess that p (W2), an unjustified belief that p (W3), the fact that one has justification for p (W4), knowledge that p (W5), a priori certainty that p (W6) and so on. It is easy to see that there is a significant difference in the degree of reliability with which assertions produced in these ways generate knowledge in hearers. In particular, W4 – W6 achieve a much higher degree of reliability than W1 – W3. Finally, assertions are produced by human agents that are free in the sense that they can produce assertions in a variety of ways, including W1 – W6, but also criticism-averse and so will, all else equal, seek to avoid producing assertions in ways that make them liable to criticism (HUMANITY).

As a result, it makes sense to regulate the production of assertions in SES-Assertion by rules that prescribe use of ways of producing assertions that are reliable in generating knowledge in hearers. So doing will lead criticism-averse asserters to favour producing assertions in ways that don't make them liable to criticism to ways that do. This will increase (or sustain) the ratio of assertions produced in reliable ways compared to assertions produced in unreliable ways and so contribute to ensuring that assertions reliably generate knowledge in hearers.

Recall that I argued that (i) when it makes sense to govern production of tokens with an e-function by a rule because it all of RELIABILITY, VARIATION and HUMANITY are satisfied; (ii) when we have independent reason to believe that production of the relevant tokens actually is governed by a certain rule; and (iii) when it can be shown that governing production of the relevant tokens by this rule does actually contribute to ensuring that they fulfil their e-function reliably, then it makes sense to govern production of these tokens by this rule in particular. In other words, we will have a rationalisation of the particular rule.

Now, note that we have just seen that, in the case of assertion, (i) holds. Moreover, §1 made a case for a relevant instance of (ii), to wit, K RULE actually governs assertion. This means that if it can be shown that the relevant instance of (iii) is true as well, i.e. that governing production of assertions in SES-Assertion by K RULE contributes to ensuring that assertions fulfil their e-function of generating knowledge in hearers reliably, we have reason to believe that it makes sense to govern production of assertion by K RULE in particular. In other words, we will have a rationalisation of K RULE.

So does operating K RULE contribute to ensuring that assertions in SES-Assertion reliably generate knowledge in hearers? As I am about to argue the answer to this question is 'yes'. To see why, note first that it is plausible that, in the vast majority of cases, a hearer acquires testimonial knowledge that p only if the speaker knows that p also. That is to say, cases in which testimony transmits knowledge predominate and cases in which testimony generates knowledge are rare exceptions. Moreover, second, absence of knowledge is often enough detectable as such by one's audience and so, if assertion is regulated by K RULE, one's audience may often enough criticise one for asserting without knowledge. Given that this is so, regulating assertion by K RULE arguably contributes to ensuring the reliability of assertion's fulfilling its e-function of generating knowledge in hearers. It does so in at least the following ways.

First, it will prevent third parties to the conversation from acquiring testimonial beliefs that fall short of knowledge. Let S be any speaker, H_1 and H_2 any two hearers and p any proposition. When S asserts that p and H_1 criticises S by saying: 'You don't know that!', H_2 will, typically at least, not acquire the belief that p based on S's testimony. Since, in the vast majority of cases in which S doesn't know that p if H_2 does acquire a belief that p based on S's testimony, H_2's belief will not qualify as knowledge, H_1's intervention contributed to improving the ratio of assertions that generate knowledge in hearers compared to assertions that don't. In this way, K RULE contributes towards ensuring the reliability of assertion's fulfilling its e-function.

Second, it will lead speakers to refrain from asserting unknown propositions in the long run. Let S be any speaker, R a range of propositions such that (i) S's beliefs about propositions in R tend to fall short of knowledge and (ii) S is disposed to assert propositions in R. If S is often enough criticised for asserting propositions in R, S's criticism-aversion will lead S to favour refraining from asserting propositions in R in the long run. Since, in the vast majority of cases in which S doesn't know

what he asserts, his hearers won't acquire knowledge based on his testimony, S's refraining from asserting propositions in R will have the effect that future hearers won't acquire testimonial beliefs from S that fall short of knowledge. Given the plausible assumption that some of these hearers would have acquired beliefs based on S's say-so, criticisms contribute to improving the ratio of assertions that generate knowledge in hearers compared to assertions that don't. K RULE contributes towards ensuring the reliability of assertion's fulfilling its e-function once again.

It comes to light that, in SES-Assertion, all of three conditions for a successful rationalisation of the rule that actually governs assertion are satisfied. If so, it not only makes sense to govern assertion by some rule or other, but it also makes sense to govern it by K RULE in particular.[20]

This completes my account of the normativity of assertion. To sum up, the practice of producing and consuming assertion is an SES. Crucially, tokens of its product, assertion, have the epistemic e-function of generating knowledge in hearers. Since e-functions give rise to evaluative norms, we can derive an evaluative norm for assertion. Moreover, there is reason to believe that, for tokens with an e-function that satisfy certain conditions (RELIABILITY, VARIATION, HUMANITY), it makes sense to govern their production by rules. Since assertion does satisfy these conditions, it makes sense to govern production of assertions by rules. Moreover, since governing production of assertions by K RULE does contribute to ensuring that assertion fulfils its epistemic e-function reliably, it makes sense to govern production of assertion by K RULE in particular.

4   Function First or Rule First?

The account I proposed, FFAA, explains the normativity of assertion in terms of its function, or, to be more precise, in terms of its epistemic e-function of generating knowledge in hearers. In this way, FFAA places function first. It differs from standard rule first accounts, which take a rule of assertion as the starting point for explaining the normativity of assertion. For instance, the rule first account that is closest to FFAA, RFAA-K, accepts a rule of assertion, K RULE, and derives an evaluative norm, THE K ACCOUNT OF GOOD ASSERTION, from it and THE K RULE ACCOUNT OF GOOD ASSERTION. I agree with champions of RFAA-K that K RULE is true. That said, FFAA offers an evaluative norm of assertion that differs from THE K ACCOUNT OF GOOD ASSERTION. Furthermore, FFAA also differs from RFAA-K in that it derives the evaluative norm from the epistemic e-function of assertion rather than from the rule. It also explains K RULE in terms of this function.

The question that arises at this stage is which of the two accounts of the normativity of assertion is preferable: the function first account I developed above or its

---

[20] Again, the rationalisation really is a prima facie one. That said, I take it that there is no alternative rule we might use to govern assertion such that our failure to use it instead of K RULE defeats the rationality of governing assertion by K RULE. If so, K RULE will also be ultima facie rationalised.

standard rule first competitor. In what follows, I will try to provide some reason to think that the balance tips in favour FFAA.

4.1 Assertions of Moorean and Lottery Propositions

Recall that RFAA-K was said to receive support from the fact that, intuitively, assertions of Moorean propositions aren't good assertions. The same goes for assertions of lottery propositions, at least when the only evidence available is the probabilistic evidence against winning. RFAA-K could easily accommodate these intuitions. Speakers don't know lottery and Moorean propositions. Since an assertion is an epistemically good one only if the speaker knows it to be true, RFAA-K predicts that assertions of lottery and Moorean propositions aren't epistemically good thus accommodating the intuition that such assertions aren't good ones.

The question that arises, then, is whether these cases allow RFAA-K to score against FFAA. Unfortunately for RFAA-K, there is reason to believe that the answer here is no as assertions of lottery and Moorean propositions come out not good on FFAA, too.

Consider lottery propositions first. When the only evidence available for a certain lottery proposition, $p - l$, is the relevant probabilistic evidence, no one is in a position to know that $p - l$. If so, an assertion of $p - l$ could not generate a knowledge that $p - l$ in hearers either. If so, the assertion is guaranteed not to have the disposition at issue in THE EVALUATIVE NORM OF ASSERTION and so is bound not to be an epistemically good assertion.

FFAA can also secure the result that assertions of Moorean propositions aren't epistemically good. Or to be more precise, it can do so given the plausible assumption that knowledge of the second conjunct of an asserted Moorean proposition— the speaker doesn't know that p—provides hearers with enough reason to distrust the first conjunct—p—that they cannot come to know that the first conjunct is true. When a speaker asserts a Moorean proposition, either hearers come to know the second conjunct or they don't. Suppose they don't. Since knowledge of a conjunction requires knowledge of each conjunct, it follows that they do not come to know the conjunction. Suppose, next, they do. By the above assumption, in that case they don't come to know the first conjunct. And since, again, knowledge of a conjunction requires knowledge of each conjunct, it follows that they do not come to know the conjunction. Either way, then, assertions of Moorean propositions could not generate knowledge of the Moorean proposition in the hearer. Any such assertion is guaranteed not to have the disposition at issue in THE EVALUATIVE NORM OF ASSERTION and so is bound not to be an epistemically good assertion.[21]

---

[21] It may also be worth noting that THE EVALUATIVE NORM OF ASSERTION predicts that assertions of false propositions are bound not to be epistemically good. Given that knowledge is factive, a false assertion that p simply could not generate knowledge that p in any hearer. If so, the assertion is guaranteed not to have the disposition at issue in THE EVALUATIVE NORM OF ASSERTION and so is bound not to be an epistemically good assertion. FFAA can thus also explain another datum Williamson adduces in his discussion of the so-called RBK account, to wit, that "we regard the false assertion itself . . . as faulty" [2000: 262].

FFAA and RFAA-K do equally well when it comes to explaining the intuition that assertions of lottery and Moorean propositions aren't good ones. There are no points to be scored by RFAA-K here.

4.2 Theoretical Motivation

I take it that one of the major attractions of FFAA is the theoretical motivation it can offer for its account of the normativity of assertion. First, the account is backed by a plausible general account of the normativity of tokens with an e-function. This account distinguishes clearly between evaluative and prescriptive norms and offers an independently motivated story of how both of these two norms are generated for items with an e-function. FFAA also offers independent arguments to the effect that assertions do have an e-function, to wit, the e-function of generating knowledge in hearers and that, as a result, the general account of the normativity of tokens with an e-function does apply to the case of assertion. In sum, according to FFAA, the normativity of assertion turns out to be but one instance of a familiar and more general kind of normativity that is associated with tokens with an e-function that distinguishes clearly between two importantly different kinds of norms and offers detailed and independently motivated accounts of both of these norms.

In contrast, champions of RFAA-K have spent a lot of time and energy on defending K RULE. However, there is little to no explicit recognition of the distinction between evaluative and prescriptive norms. Unsurprisingly, then, it is simply assumed, and often only implicitly, that a good assertion is an assertion that satisfies K RULE, an idea that is captured in RFAA's second key thesis THE K RULE ACCOUNT OF GOOD ASSERTION. This means that champions of RFAA-K have simply assumed one out of two key theses of their account, rather than supported it by independent argument. Once we are clear on the fact evaluative and prescriptive norms are different kinds of norms and once we have seen that they need not be related in the way champions of RFAA-K have assumed, it should also be clear that it won't do to just assume that the conditions for satisfying the evaluative norm coincide with the conditions required for satisfying the prescriptive norm. If they do, this point affords independent argument. Since champions of RFAA-K have not produced any such argument, there is an important lacuna in their account.

In fact, things may be worse than this. After all, FFAA not only clearly distinguishes between the evaluative and the prescriptive norm of assertion. FFAA also offers independent support for its proposed evaluative norm which derives from the general account of the normativity of tokens with an e-function and the argument that assertion has an e-function. Given that this is so, if it can be shown that FFAA's and RFAA-K's evaluative norms come apart, this will serve to exert pressure on RFAA-K. After all, there is independent theoretical reason to believe that FFAA's norm is true, whereas there is no such reason to believe that RFAA-K's norm is true. If so, there is some theoretical reason to believe that, in cases in which the two come apart, FFAA's verdict is the correct one. In what follows, I will look at one case that fits the bill and argue that it serves to establish a further advantage for FFAA.

4.3 Cases of Selfless Assertion: Theoretical Considerations

Jennifer Lackey has offered a kind of case that involve what she calls selfless assertion. Consider the following version of what is perhaps the most famous example of such a case:

> *Creationist Teacher.* Stella is a teacher who, whilst herself being a creationist, recognises that the scientific evidence strongly supports evolutionary theory. Since she takes it to be her duty as a teacher to present the view that is supported by scientific evidence, she asserts to her students various truths of evolutionary theory, including that modern-day Homo sapiens evolved from Homo erectus.[22]

Lackey takes this case to constitute a counterexample to K RULE. Being a creationist, Stella does not believe and hence does not know that Homo sapiens evolved from Homo erectus. Hence her assertion violates K RULE. At the same time, intuitively, her assertion is good. So, K RULE is in trouble. With Lackey's assessment of the case on the table, I'd like to set it aside for now. I will return to it in due course.

What I would like to focus on instead is what FFAA and RFAA-K have to say about this case. Let's start with RFAA-K. Recall RFAA-K's key evaluative norm of assertion, THE K ACCOUNT OF GOOD ASSERTION: one's assertion that p is (epistemically) good if and only if one knows that p. Since it is undeniable that Stella does not know that Homo sapiens evolved from Homo erectus, THE K ACCOUNT OF GOOD ASSERTION delivers the result that Stella's assertion is not an epistemically good one.

Let's now take a look at FFAA. Recall that FFAA's key evaluative norm of assertion is THE EVALUATIVE NORM OF ASSERTION: one's assertion that p is (epistemically) good if and only if it has the disposition to generate knowledge in hearers by functioning normally when in normal conditions. Crucially, there is reason to believe that Stella's assertion does have this disposition. After all, when Stella's students acquire the corresponding beliefs about the evolution of Homo sapiens, their beliefs will qualify as knowledge. That is to say, in this case, Stella's assertion does generate the relevant knowledge in her students. If so, in this case, there is reason to believe that Stella's assertion manifests the disposition to generate knowledge in students. What's more, it does so in just the same way as the assertions of a speaker who knows these truths him or herself in the same conditions. If so, it is plausible that her assertion manifests the disposition by functioning normally in normal conditions. Since her assertion can manifest only dispositions it also has, there is reason to believe that Stella's assertions does have the disposition to generate knowledge in hearers by functioning normally when in normal conditions. According to THE EVALUATIVE NORM OF ASSERTION, her assertion is an epistemically good one.

It comes to light that FFAA and RFAA-K do come apart in the sense that they make diverging predictions about *Creationist Teacher*: FFAA predicts that Stella's assertion is good, whereas, according to RFAA-K, it isn't. Since as we saw above,

---

[22] [Lackey 2008: 48]. Note that Douven [2006: 461] also offers a similar kind of case although he does not explicitly use it to target K RULE.

FFAA's evaluative norm is independently motivated, whereas RFAA-K's is simply assumed to be true, there is independent theoretical reason to think that FFAA's prediction is the right one here.

Moreover, according to FFAA, *Creationist Teacher* is but one instance of a more general phenomenon that arises for tokens with an e-function the production of which is also governed by a prescriptive norm. There are other non-epistemic cases in which token products with an e-function satisfy the relevant evaluative norm even though their producers violated the rule governing their production. By way of illustration, let's return once more to SES-Espresso. Note that a cup of espresso may satisfy the evaluative culinary norm of being a good cup of espresso: it does have the disposition to generate pleasant gustatory experiences in consumers by functioning normally when in normal conditions. At the same time, the barista who made it may have done just about everything wrong when making it. That is to say, he may have broken just about any rule of espresso making there is, including the rule operative at the local coffee shop. What we have is a case of a cup of espresso that is good despite the fact that it was produced in violation of the rule for making good cups of espresso. The fact that, according to FFAA, *Creationist Teacher* turns out to be but another instance of a general phenomenon adds to making FFAA's analysis a theoretically appealing one.

4.4 Cases of Selfless Assertion: Intuitive Considerations

With these theoretical considerations in play, I would now like to return to Lackey's analysis of cases selfless assertion. Recall that I said that Lackey takes these cases to constitute counterexamples to K RULE as the speakers' assertions are good and yet they violate K RULE. It is worth noting that cases of selfless assertion are not taken to be just any old problem for K RULE. Rather, they are thought to be a particularly difficult problem for the view. The reason for this is that other apparent problem cases—such as intuitively good assertions in Gettier cases and in cases in which the speaker has a justified false belief—can be explained by champions of K RULE in terms of blamelessness. The thought here is that, while the assertions violated K RULE, the speakers are blameless for making these assertion. After all, they reasonably believe that they are complying with K RULE. Clearly, however, this will not work for cases of selfless assertion. Selfless asserters such as Stella may be well aware that they are violating K RULE. As a result, the intuitive goodness of selfless assertions cannot be explained in terms of blamelessness. That's why, cases of selfless assertions pose a particularly hard problem for K RULE.[23]

Do cases of selfless assertion refute K RULE? Unsurprisingly, I think the answer to this question is no. The reason for this already surfaced earlier, in §1: if a student of Stella's were to know that Stella does not believe that Homo sapiens evolved from Homo erectus, she may prima facie legitimately criticise Stella for her

---

[23] Note, furthermore, that the perhaps most obvious alternative response to the case, to wit, that selfless asserters don't assert in their own names but in accordance with the requirements of a social role, does not appear to work. [Lackey 2008] and [Douven and Kelp 2012] both offer examples of cases of selfless assertion in which the speaker does not occupy a certain social role when asserting.

assertion, for instance by saying: 'You don't believe that yourself!' Moreover, this criticism would call for or make appropriate an excuse, explanation etc. on the part of Stella. A rebuttal of the criticism as misplaced or irrelevant would not be prima facie legitimate. For instance, 'So what?' is not a prima facie legitimate response. As we have already seen, these considerations suggest that K RULE is in full force here.

But then what exactly is going on in Lackey's argument? Let's first take another look at it. There are two important claims here. First, Stella's assertion violates K RULE. This much is hard to deny. Second, intuitively, Stella's assertion is good.

Now, we have just seen that the case does not serve to refute K RULE. Even so, we still need an account of the intuitive goodness of Stella's assertion. It is here that FFAA can again score points. After all, according to FFAA, Stella's assertion is actually an epistemically good one. In this way, FFAA can accommodate the intuition that Stella's assertion is good in a straightforward way. As a result, FFAA can offer an appealing account of a kind of case that has long been a thorn is the side of K RULE: FFAA can allow (i) that K RULE is violated in this case, whilst (ii) plausibly resisting the idea that this provides convincing reason to think that K RULE is not in force. Finally, (iii) FFAA can also accommodate the intuition that assertions like Stella's are epistemically good in a straightforward way.

Of course, this kind of account is not available to champions of RFAA-K. After all, according to RFAA-K, the conditions for epistemically good assertion are satisfied if and only if the speaker satisfies K RULE. Given that selfless assertions clearly don't satisfy K RULE, according to RFAA-K, these assertions are bound not to be epistemically good. FFAA's appealing account of cases of selfless assertion is not available to RFAA-K. What transpires, then, is that cases of selfless assertion serve to give FFAA an advantage over RFAA-K not only when it comes to theoretical considerations, but also when it comes to intuitive ones.

Before moving on, I would like to briefly offer a diagnosis of what went wrong in RFAA-K. The discussion of theoretical considerations indicated that one of RFAA-K's key theses, THE K RULE ACCOUNT OF GOOD ASSERTION, is simply assumed to be true rather than independently motivated. In addition, the discussion of intuitive considerations showed that RFAA-K's evaluative norm, THE K ACCOUNT OF GOOD ASSERTION, comes under pressure by cases of selfless assertion. But now recall that THE K ACCOUNT OF GOOD ASSERTION is entailed by K RULE and THE K RULE ACCOUNT OF GOOD ASSERTION. Since K RULE is independently plausible and, as we have just seen, not impugned by *Creationist Teacher*, this means that the intuitive pressure exerted by cases of selfless assertion on THE K ACCOUNT OF GOOD ASSERTION transmits to THE K RULE ACCOUNT OF GOOD ASSERTION. What comes to light, then, is that THE K RULE ACCOUNT OF GOOD ASSERTION is not only wanting independent support, but also comes under intuitive pressure. It is thus tempting to think that RFAA-K fails because THE K RULE ACCOUNT OF GOOD ASSERTION is false.

4.5   Rationalisations of K Rule

So far we have seen that FFAA compares favourably with RFAA-K in the following respects: it enjoys a higher quality theoretical motivation and it can offer a better account of a particularly difficult kind of case for K RULE, that of selfless assertion, on both a theoretical and an intuitive level. I will now argue that FFAA can score points against RFAA-K on yet another count, this time relating to the prescriptive norm of assertion.

Recall that RFAA-K and FFAA agree on the prescriptive norm of assertion. Both accept that K RULE is true. The mere fact that we have an answer to the question of whether K RULE governs assertion does of course not mean that we have an answer to the question of whether it makes sense for assertion to be governed by K RULE. The question of a rationalisation of K RULE remains to be addressed.

Now notice that I have argued above that FFAA can offer an appealing answer to this question. Here it is again, in thumbnail version. (i) For tokens with an e-function such that RELIABILITY, VARIATION and HUMANITY are satisfied, it makes sense to govern their production by a rule. What's more, when (ii) there is independent reason to believe that there actually is a particular rule governing productions of such tokens with an e-function, it makes sense to govern production of such tokens by this particular rule if (iii) it can be argued that operating the particular rule contributes to ensuring that tokens do fulfil their e-function reliably. This means that if it can be shown that (i) - (iii) hold, it makes sense to govern production of the relevant tokens by the particular rule mentioned in (ii). We will have a rationalisation of this rule. Moreover, I have argued at length that all of (i) – (iii) are satisfied in the case of assertion, i.e. that (i) assertions have the e-function of generating knowledge in hearers and that RELIABILITY, VARIATION and HUMANITY are satisfied, (ii) that K RULE actually governs production of assertion and (iii) that governing assertion by K RULE does actually contribute to ensuring that assertions fulfils their e-function reliably. In this way, FFAA can offer a rationalisation of K RULE, and an attractive one at that.

Can RFAA-K do equally well on this front? As a first observation, note that the question of rationalisation is not often addressed by contributors to the debate on rules of assertion. However, one answer can be reconstructed from Williamson's account of the nature of assertion. Another one, advocated by Kent Bach and Frank Hindriks, takes on the question more directly.[24] In what follows, I will take a closer look at both of these answers, starting with Williamson's.

WILLIAMSON

We can make some progress towards a Williamsonian treatment of answer to the question of why assertion should be governed by K RULE by looking at his account of the nature of assertion. According to Williamson, K RULE is *constitutive* of

---

[24] Yet another noteworthy rationalisation of a rule of assertion is can be found in [Douven 2006]. Since Douven's argument aims to rationalise the justification rule of assertion (see n.1), it will not be necessary to discuss it here.

assertion in much the same way in which rules of a game are constitutive of games [Williamson 2000: 238].[25] Given that this is so, we must not expect there to be a very deep answer to our question. In order to see this, consider the case of games. To make things a bit more concrete, let's take chess as an example. Why should chess be governed by one set of rules rather than another? For instance, why should chess be governed by a rule according to which the pawn can move two squares forward if it has not yet been moved, but only one square if it has already been moved (R1), rather than by a rule according to which it can always only move one square (R2)? One would not expect questions like this one to have deep answers. We might easily have come to play a game that is governed by R2 rather than R1. Of course, in that case, we would not have played chess, but some different game. While in this particular case, the game we might have played is still quite similar to chess, there is not even special reason to believe that we should not have played an entirely different game instead. That certain games, including chess, became popular, whilst others didn't, is plausibly an historical contingency. Given that rules of assertion are constitutive of assertion in an analogous manner, there is reason to believe that the same goes for assertion. We now have ingredients for the Williamsonian treatment of the rationalisation question: there is independent reason to think that we should not expect there to be a deep answer to this question in the first place. And, of course, given that this is so, the fact that no rationalisation has been given constitutes no disadvantage for RFAA-K.

At the same time, it seems somewhat less plausible that, in the case of rules of assertion, it is a mere historical contingency that a speech act governed by K RULE, say, came to be widely used. Supposing that assertion really is governed by K RULE, it is somewhat less plausible to think that we might as well have ended up using a speech act that is governed by different rule instead, especially when this rule is quite different than the rule that actually governs assertion. For instance, it is rather implausible that we might as well have ended up using a speech act governed by the following rule of assertion:

THE PLEASURE RULE OF ASSERTION. One must: assert *p* only if *p* pleases one.

It seems that there should be a deeper answer to the question of why assertion should be governed by K RULE. So long as no answer has been given by champions of RFAA-K, they are at a disadvantage.

What's more, there is even reason to worry that the above line of reasoning backfires. If the analogy in conjunction with the fact that there is no deep answer to the rationalisation question in the case of games provides reason to believe that we must not expect there to be a deep answer to the rationalisation question in the case of assertion, then the fact that there should be some such answer in the case of assertion means trouble for the Williamsonian. After all, these considerations suggest that champions of the Williamson line have not only failed to provide a answer but also that there might just be no satisfactory answer forthcoming here.

---

[25] In fact, according to Williamson, K RULE is the unique constitutive rule of assertion and assertion is the only speech act whose unique constitutive rule is K RULE [Williamson 2000: 241].

If so, the disadvantage is not one that we may expect to be remedied by further research.[26],[27]

BACH AND HINDRIKS

Bach [2008, 2010] has proposed a different rationale for K RULE. Roughly, the idea is that K RULE can be derived from the idea that assertion is the linguistic expression of belief, which gives rise to a belief rule on assertion, and a knowledge rule on belief. This idea has subsequently been developed in more detail by Frank Hindriks. Starting from the idea that to assert that p is to utter a sentence that means that p and thereby expresses the belief that p (1), Hindriks argues that, in situations of normal trust, which obtain unless it is permissible to lie, assertion is governed by a sincerity rule to the effect that one must: express the belief that p only if one believes that p (2). (1) and (2) imply that, in situations of normal trust, one must assert that p only if one believes that p (3), which is a restricted belief rule on assertion. In conjunction with the knowledge rule of belief[28], according to which one must: believe that p only if one knows that p (4), we get the result that, in situations of normal trust, one must assert that p only if one knows that p (5). And this, of course, is a restricted version of K RULE [Hindriks 2007: 403].

Now, there are a number of points on which one might disagree with Bach and Hindriks.[29] For the purposes of this paper, I will focus on one: rules are not transitive in the sense that if there is a rule to the effect that one must: ϕ only if condition C obtains, and a rule to the effect that one must: ψ (where ψ-ing entails that condition C obtains) only if condition D obtains, then there is also a rule to the

---

[26] Of course, champions of the Williamson line might maintain that the analogy between rules of assertion and rules of games does not extend to the rationalisation question in the following sense: the fact that we must not expect there to be a deep answer to the rationalisation question in the case of games does *not* provide reason to think that there is no deep such answer in the case of assertion either. Note, however, that, in that case, Williamsonians will not only shoulder the burden of offering a satisfactory answer to the rationalisation question. In addition, they will also need to argue that the analogy does not extend to the rationalisation question. After all, the onus is on them to show that the analogy between rules of assertion and games does not extent to just the point where it becomes potentially problematic for them.

[27] In fact, as I argue elsewhere [Author 2014b] in more detail, there is excellent reason to believe that speaker rules governing assertion are not constitutive rules. The above argument that champions of constitutivity won't be able to offer the right kind of explanation of why assertion should be governed by one speaker rule rather than another confirms this result.

[28] For more on the knowledge rule of belief, see e.g. [Adler 2002, Gibbons 2013, Sosa 2011, Williamson 2000].

[29] Brian Ball [2014a] has charged Hindriks with an equivocation as he takes the belief rule of assertion to be a moral norm and the knowledge rule of belief an epistemic norm. For more on this see [Hindriks and Kooi 2014] and [Ball 2014b]. I also have reservations about the knowledge rule of belief. The reason for this is that the knowledge rule of belief (it is (epistemically) permissible for one to believe that p only if one knows that p) in conjunction with the apparently innocuous thesis that if one justifiably believes that p then it is (epistemically) permissible to believe that p, entails that one justifiably believes that p only if one knows that p. However this is highly counterintuitive. Moreover, as I argue elsewhere in more detail [Author 2011, 2015b], there is excellent reason to think that it is false.

effect that one must: ϕ only if condition D obtains.

To see this, consider first the following counterexample due to Clayton Littlejohn. Suppose there is a first rule to the effect that one must: apologise for breaking the neighbour's window only if one breaks the neighbour's window and a second rule to the effect that one must: break the neighbour's window only if the neighbour has given one permission to break it. Even so, there is no rule to the effect that one must: apologise for breaking the neighbour's window only if the neighbour has given one permission to break it. If so, rules are not transitive in the above sense.[30]

Note that Littlejohn's counterexample involves a contrary-to-duty rule. That said, there are other counterexamples to transitivity that do not exploit cases involving contrary-to-duty rules. Here is one: suppose that there is a rule to the effect that one must: admit a student to the university only if he has been issued a high school diploma. Suppose that there is another rule to the effect that one must: issue a high school diploma to a student only if he has paid his high school tuition fees. It simply does not follow that there is a rule to the effect that one must: admit a student to the university only if he has paid his high school tuition fees. If such a rule exist, it must have been introduced separately.

By way of evidence, note that if we were to criticise someone for admitting a student to university even though he hasn't paid his high school tuition fees, for instance, by saying: 'The student has not paid his high school tuition fees!', our criticism would be prima facie illegitimate, even if we know that he hasn't paid and can provide reason in support of it. The agent criticised may prima facie legitimately rebut our criticism as misplaced, perhaps by telling us to direct it at the person who issued the diploma. (And even if the person who grants admission is the same person as the one who issued the diploma, that person may point out that the problem is a different one, pertaining to the issuing of the high school diploma rather than the admission to university.) If we want it to be the case that the person admitting the student can be prima facie legitimately criticised for admitting him despite not having paid his high school tuition fees, we need to introduce a separate rule to this effect. This is exactly as LEGITIMATE CRITICISMS, ILLEGITIMATE CRITICISMS and LEGITIMATE REBUTTALS would have it. These considerations suggest that the prospects of deriving K RULE from the belief rule of assertion and the knowledge rule of belief are dim.

In sum, we have seen that FFAA can offer an attractive account of why assertion should be governed by K RULE. In contrast, it turns out that extant rationalisations of K RULE in the literature that are congenial to RFAA-K remain unsatisfactory. FFAA thus compares favourably with RFAA-K in yet another respect.

Conclusion

Standard accounts of the normativity of assertion are rule first accounts. They accept that assertion is governed by THE C RULE OF ASSERTION and that whether or not

---

[30] Littlejohn gives this example on his blog Think Tonk (`http://claytonlittlejohn.blogspot.be/2013/04/sincerity-assertion-and-case-for-common.html`).

a particular assertion is a good one depends on whether the speaker satisfies THE C RULE OF ASSERTION. The perhaps most prominent rule first account of assertion is RFAA-K, which identifies the crucial property C with knowledge.

This paper has argued that RFAA-K is right on one count: assertion is governed by K RULE. This is confirmed by a number of independently plausible principles for performance types that are/aren't governed by a rule, to wit, LEGITIMATE CRITICISMS, ILLEGITIMATE CRITICISMS and LEGITIMATE REBUTTALS.

The main ambition of the paper has been to argue that, besides K RULE, RFAA-K is not all that attractive. A more promising alternative account of the normativity of assertion is available. I have developed and defended this account, FFAA, in some detail. In particular, I have offered an independently plausible account of e-functions. With this account in play I have argued that e-functions give rise to evaluative norms and that, given certain conditions, they also serve to rationalise prescriptive norms governing production of e-functional items. With this general account of e-functions and their normative import in play, I have argued that assertion has the epistemic e-function of generating knowledge in hearers and showed that, as a result, we get an evaluative norm of assertion as well as a rationalisation of K RULE.

Finally, I have argued that FFAA compares favourably with RFAA-K. First, FFAA can offer an equally plausible account of why assertions of lottery and Moorean propositions aren't good ones, which means that there is no advantage to be gained for RFAA-K here. Second, FFAA offers an independently motivated account of the normativity of assertion that distinguishes clearly between evaluative and prescriptive norms. In contrast, champions of RFAA-K do not distinguish explicitly between these two kinds of norms. As a result, their evaluative norm remains without independent support. Third, there is then theoretical reason to favour FFAA's account in cases in which FFAA and RFAA-K make diverging predictions such as cases of selfless assertion. Fourth, this theoretical reason is confirmed by the intuitive verdicts in such cases. FFAA can offer an appealing account of the intuition of the goodness of selfless assertions whilst holding on to K RULE, thus offering an appealing solution to a particularly difficult problem for champions of K RULE. Finally, fifth, the rationalisations of K RULE provided by champions of RFAA-K remain unsatisfactory, whereas FFAA can offer an appealing account here.

In view of these results, I submit, FFAA is a promising account of the normativity of assertion that deserves to be taken seriously.

References

Adler, J. 2002. *Belief's Own Ethics*. MIT Press, Cambridge/MA.

Author 2011. Author's work.

Author 2014a. Author's work.

Author 2014b. Author's work.

Author 2015a. Author's work.

Author 2015b. Author's work.

Bach, K. 2008. Applying pragmatics to epistemology. *Philosophical Issues*, 18:68–88.

Bach, K. 2010. Knowledge in and out of context. In Campbell, J., O'Rourke, M., and Silverstein, H., editors, *Knowledge and Skepticism*. MIT Press, Cambridge/MA.

Ball, B. 2014a. Deriving the norm of assertion (forthcoming). *Journal of Philosophical Research*.

Ball, B. 2014b. Response to Hindriks and Kooi (forthcoming). *Journal of Philosophical Research*.

Bigelow, J. and Pargetter, R. 1987. Functions. *Journal of Philosophy*, 86:181–96.

Buller, D. 1998. Etiological theories of function: a geographical survey. *Biology and Philosophy*, 13:505–527.

DeRose, K. 2002. Assertion, knowledge, and context. *Philosophical Review*, 111:167–203.

Douven, I. 2006. Assertion, knowledge, and rational credibility. *Philosophical Review*, 115:449–485.

Douven, I. 2009. Assertion, Moore and Bayes. *Philosophical Studies*, 144:361–75.

Douven, I. and Kelp, C. 2012. In defense of the rational credibility account: a reply to Casalegno. *dialectica*, 66:289–97.

Geach, P. 1956. Good and evil. *Analysis*, 17:32–42.

Gerken, M. 2012. Discursive justification and skepticism. *Synthese*, 189:373–394.

Gibbons, J. 2013. *The Norm of Belief*. Oxford University Press, Oxford.

Godfrey-Smith, P. 1993. Functions: consensus without unity. *Pacific Philosophical Quarterly*, 74:196–208.

Godfrey-Smith, P. 1994. A modern history theory of functions. *Noûs*, 28:344–62.

Goldberg, S. 2015. *Assertion: On the Philosophical Significance of Assertoric Speech*. Oxford University Press, Oxford.

Graham, P. 2010. Testimonial entitlement and the function of comprehension. In Haddock, A., Millar, A., and Pritchard, D., editors, *Social Epistemology*. Oxford University Press, New York.

Graham, P. 2012. Epistemic entitlement. *Noûs*, 46:449–82.

Graham, P. 2014a. The function of perception. In Fairweather, A., editor, *Virtue Epistemology Naturalized. Bridges between Virtue Epistemology and Philosophy of Science*. Springer, Dordrecht.

Graham, P. 2014b. Functions, warrant, history. In Fairweather, A. and Flanagan, O., editors, *Naturalizing Epistemic Virtue*. Cambridge University Press, Cambridge.

Griffith, P. 1993. Functional analysis and proper functions. *British Journal for the Philosophy of Science*, 44:409–22.

Hindriks, F. 2007. The status of the knowledge account of assertion. *Linguistics and Philosophy*, 30:393–406.

Hindriks, F. and Kooi, B. 2014. Reaffirming the status of the knowledge account of assertion (forthcoming). *Journal of Philosophical Research*.

Kitcher, P. 1994. Function and design. *Midwest Studies in Philosophy*, 18:379–97.

Kvanvig, J. 2009. The value of understanding. In Haddock, A., Millar, A., and Pritchard, D., editors, *Epistemic Value*. Oxford University Press, Oxford.

Kvanvig, J. 2010. Norms of assertion. In Brown, J. and Cappelen, H., editors, *Assertion*. Oxford University Press, Oxford.

Lackey, J. 2007. Norms of assertion. *Noûs*, 41:594–626.

Lackey, J. 2008. *Learning from Words: Testimony as a Source of Knowledge*. Oxford University Press, Oxford.

McHugh, C. 2012. The truth norm of belief. *Pacific Philosophical Quarterly*, 93:8–30.

McKinnon, R. 2013. The supportive reasons norm of assertion. *American Philosophical Quarterly*, 50:121–35.

McKinnon, R. 2015. *Norms of Assertion: Truth, Lies, and Warrant*. Palgrave Macmillan, Basingstoke.

Millikan, R. 1984. *Language, Thought, and Other Biological Categories*. MIT Press, Cambridge/MA.

Millikan, R. 2004. *Varieties of Meaning*. MIT Press, Cambridge/MA.

Neander, K. 1991a. Functions as selected effects: the conceptual analyst's defence. *Philosophy of Science*, 58:168–84.

Neander, K. 1991b. The teleological notion of 'function'. *Australasian Journal of Philosophy*, 69:454–68.

Simion, M. 2015. Assertion: knowledge is enough (forthcoming). *Synthes*.

Slote, M. 1979. Assertion and belief. In Dancy, J., editor, *Papers on Language and Logic*. Keele University Press, Keele.

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

Turri, J. 2011. The express knowledge account of assertino. *Australasian Journal of Philosophy*, 89:37–45.

Turri, J. 2014. Knowledge and suberogatory assertion. *Philosophical Studies*, 167:557–67.

Unger, P. 1975. *Ignorance: A Case for Scepticism*. Clarendon, Oxford.

Walsh, D. and Ariew, A. 1996. A taxonomy of functions. *Canadian Journal of Philosophy*, 26:493–514.

Weiner, M. 2005. Must we know what we say? *Philosophical Review*, 114:227–251.

Whiting, D. 2013. Stick to the facts: on the norms of assertion. *Erkenntnis*, 78:847–67.

Williamson, T. 1996. Knowing and asserting. *Philosophical Review*, 105:489–523.

Williamson, T. 2000. *Knowledge and Its Limits*. Oxford University Press, Oxford.