



Knowledge and reasoning

Mona Simion¹

Received: 1 April 2020 / Accepted: 4 June 2021 / Published online: 21 June 2021
© The Author(s) 2021

Abstract

This paper develops a novel, functionalist, unified account of the epistemic normativity of reasoning. On this view, epistemic norms drop out of epistemic functions. I argue that practical reasoning serves a prudential function of generating prudentially permissible action, and the epistemic function of generating knowledge of what one ought to do. This picture, if right, goes a long way towards normatively divorcing action and practical reasoning. At the same time, it unifies reasoning epistemically: practical and theoretical reasoning will turn out to be governed by the same epistemic norm—knowledge—in virtue of serving the same epistemic function: generating knowledge of the conclusion.

Keywords Knowledge norm · Practical reasoning · Theoretical reasoning · Action

1 Introduction

The specialist literature tends, with very few exceptions,¹ to lump discussions of the epistemic norm of action together with discussions of the epistemic norm of practical reasoning.² In contrast, discussions of the epistemic normativity of theoretical

¹ See e.g. Fantl and McGrath (2009, 2012), Simion (2018, 2020).

² This paper is concerned with the debate concerning the epistemic status that a premise needs to enjoy in order to be permissibly employed in practical/theoretical reasoning (e.g. Brown, 2010, 2012; Douven, 2008; Hawthorne & Stanley, 2008; Littlejohn, 2014; Neta, 2009; Reed, 2012; Simion, 2018; Stanley, 2005; Weatherson, 2012; Williamson, 2000). There is a parallel debate concerning the normativity of reasoning that discusses the normativity governing the transition from one premise to another/conclusion, which does, to a large extent, discuss reasoning *simpliciter*—be it theoretical or practical (e.g. Broome, 2013; Dancy, 2004, 2018; McHugh and Way, 2008).

This article belongs to the topical collection “Knowledge and Decision”, edited by Moritz Schulz, Roman Heil, Jakob Koscholke, and Patricia Rich.

✉ Mona Simion
mona.simion@glasgow.ac.uk

¹ COGITO Epistemology Research Centre, University of Glasgow, 69 Oakfield Avenue, Glasgow G12 8LP, UK

reasoning tend to be carried separately from discussions of the normativity of practical reasoning. This paper develops a novel, integrated account of the normativity of reasoning that breaks with this norm.

On the view proposed here, all epistemic norms are generated by associated epistemic functions, in the same way in which biological norms are generated by biological functions. Since epistemic norms are associated with epistemic functions, epistemic norms will only govern particular action types that have distinctively epistemic functions, not action in general. Compatibly, action in general is governed by prudential norms with epistemic content, i.e. regulating the amount of epistemic support required for prudentially appropriate acting.

While this view predicts that there is no epistemic norm for action in general, the question as to whether there is an epistemic norm for practical reasoning remains open. I argue that practical reasoning serves (1) in line with action, moral and prudential functions of generating morally/prudentially permissible action,³ but also (2) the epistemic function of generating knowledge of what one ought to do. In virtue of this latter function, practical reasoning will be governed by a corresponding epistemic norm. This picture, if right, sharply separates action and practical reasoning normatively. At the same time, it provides us with an exciting opportunity to unify reasoning⁴: practical and theoretical reasoning will turn out to be governed by the same epistemic norm—knowledge—in virtue of serving the same epistemic function: generating knowledge of the conclusion.

Here is the game plan: Sect. 2 outlines the functionalist normative framework I favour. In Sect. 3, I defend my nihilism about the epistemic normativity of action. Section 4 discusses practical and theoretical reasoning: I argue for a knowledge generating function of reasoning and for the corresponding knowledge norms dropping out of this function. In the last section I conclude.

2 Knowledge functions

2.1 Normative functionalism

This paper develops a view on which reasoning is a move in our epistemic practice of inquiry, just like moving the bishop diagonally is a move in our ludic practice of chess. This paper conjectures that, in line with the normativity of moves in practices in general, we should expect the epistemic normativity of reasoning to drop out of its status in the practice of inquiry.

³ I don't take a stance here on the question of whether practical reasoning's main function is prudential, moral, or something more highbrow like eudaimonia, flourishing etc. I will assume a prudential function for my purposes here, but everything I say will apply, *mutatis mutandis*, to any of the candidates on the market. See (Wallace, 2020) for an excellent overview.

⁴ See also e.g. Littlejohn (2014), Wedgwood (2002).

I want to get one thing out of the way right from the start: when one discusses epistemic norms for reasoning, there are two things one can be talking about: transition epistemic norms—governing how I should handle the premises and move from them to the conclusion—and epistemic norms pertaining to what premises I am permitted to rely on to begin with. The issue this paper is concerned with (in line with the ‘epistemic norms of assertion, action, and belief’ debates, e.g. (Benton, 2014)) is the latter: the question I am trying to answer here is: ‘what properties need a proposition enjoy in order to be epistemically permissibly employed in reasoning?’ and the answer I defend is: it needs to be known.

In several places (e.g. Simion, 2018, 2019, Forthcoming), I have defended a normative picture on which all norms are sourced in functions.⁵ In traits, artefacts and practices alike, functions generate norms. There is such a thing as a properly functioning heart, a properly functioning can opener and a proper way to make coffee. If that is so, when we are interested in a particular type of norm governing a particular type of practice, it is helpful to first identify its function.⁶

On the etiological theory of functions,⁷ functions turn on histories that explain why the item exists or operates the way it does. Take my heart; plausibly, tokens of the type pumped blood in my ancestors. This was beneficial for my ancestors’ survival, which explains why tokens of the type continue to exist. As a result, my heart acquired the etiological function (henceforth also e-function) of pumping blood. Acquiring an etiological function is a success story: traits, artefacts and actions get etiological functions of a particular type by producing the relevant type of benefit. My heart acquired a biological etiological function by generating biological benefit. Through a positive feedback mechanism—the heart pumped blood, which kept the organism alive, which, in turn, insured the continuous existence of the heart—our hearts acquired the etiological function of pumping blood.

While etiology does require some history, it is important to note that it does not require an awful lot of it; there are several ways to cash out the etiological requirement that do not presuppose directional selection, i.e. selection over generations.⁸ A trait can also acquire a particular function by ongoing, maintenance selection, or

⁵ See also, e.g., Graham (2012), Millikan (1984), Sullivan-Bisset (2017).

⁶ I am not arguing for the functionalist picture here (I have done so elsewhere (Simion, 2019, 2020, 2021) at considerable length), I assume it, develop the corresponding views, and show that they nicely explain the data that competing views have been struggling with. That being said, I do not make the corresponding descriptive claim—i.e. the claim that reasoning is a functional kind, i.e. a kind that only affords a functionalist analysis, and not a dismantling analysis. In a nutshell: for my purposes here, I only subscribe to normative functionalism, not metaphysical functionalism.

⁷ Defended by people like e.g. David J. Buller (1998), Ruth Millikan (1984), Karen Neander (1991). The etiological theory of functions is, by far, the most widely endorsed view in the literature. Its main competitor is the ‘systemic’ theory of functions, notably defended in (Cummins, 1975). Systemic functions describe how something works or operates—what it does—as a part of a larger system. Functions, in this sense, are the causal role capacities of parts that contribute to some capacity of the containing system. Systemic functions are widely believed to lack normative import, which is what explains, to a large extent, both the popularity of the competing, etiological account and why the latter is thought to be much better suited for applications to normative domains like epistemology.

⁸ See (Buller, 1998) for an excellent overview of etiological theories of function.

through a learning process, or even by the metabolic activity of the organism itself. What it all amounts to is explaining the existence/continuous existence of a trait through a longer or shorter history of positive feedback (Graham, 2014, p. 35).

Before moving on, one more important thing to be discussed is the typing of the relevant e-functions; after all, functions can be of different sorts: there are biological functions, aesthetic functions, social functions, etc. The account I favour⁹ departs from the way in which the e-functionalist picture has been traditionally applied to epistemological matters (and outside biology in general) by people like Peter Graham (2012) and Ruth Millikan (1984) in one crucial way: Graham and Millikan take biological benefit to correspond to the acquisition of all types of e-functions. In contrast, the account I favour takes functions to be typed by the corresponding benefit. As such, on my account, if a trait produces a benefit of type B in a system, the function thereby acquired will be a function of type B. If it is biological benefit that is at stake in function acquisition, what we get is a biological function. On this view, the heart's function to pump blood is a biological function in virtue of the fact that the produced benefit is also biological—i.e., survival. The function of music is an aesthetic function in virtue of the fact that the produced benefit is an aesthetic benefit. And so on. Now, of course, aesthetic benefit might, and often will, also result in biological benefit. What is important to keep in mind, though, is that the benefit that is *essential* to aesthetic function acquisition is the aesthetic one. The fact that biological benefit is also associated with the latter is, at best, a contingent matter of fact. Here is, then, the full etiological account to be employed throughout this paper:

E-FUNCTION: A token of type T has the e-function of type B of producing effect E in system S iff (1) tokens of T produced E in the past, (2) producing E resulted in benefit of type B in S or S's ancestors and (3) producing E's having B-benefitted S or S's ancestors contributes to the explanation of why T continues to exist in S.

Now, with the full account in play, note that the etiological account is an account of functions as purposes: by being selected for it, our hearts have acquired the purpose of pumping blood in our organisms. Reaching that purpose—i.e., successfully pumping blood—will amount to function fulfilment. But purposes will also come with associated norms prescribing the right way to proceed in order to reliably enough¹⁰ reach the corresponding purpose in normal conditions. Because its function contributes to the explanation of its very existence, the trait in question *ought* to perform in a way that is associated with likely enough function fulfilment. Now, according to the etiological theory of functions, this is but the way in which the trait functioned back in the day when it acquired its function. Your heart will pump blood in normal conditions, i.e., conditions similar to those in which it was selected, when functioning normally, that is, when functioning in the way in which it was functioning when it was selected for its beneficial effects. Plausibly, in normal conditions, a normally functioning heart will fulfil its function of pumping blood in

⁹ I defended this picture at length in several places, e.g. Simion (2019, 2020, 2021).

¹⁰ How reliable is reliably enough? The threshold varies with the type of trait and evolutionary needs. The heart needs to be very reliable for species continuous existence; in contrast, sperms only need to fulfill their function very rarely.

your system by beating at a particular rate. According to the etiological theory, then, normal functioning is proper functioning: a heart functions in the way it ought to function (i.e., by the norm) when it functions in the way it did back in the day when it acquired its function: when it beats at that particular rate. When functioning normally—whether in normal conditions or not—your heart will function properly, i.e. it will meet the norm constitutively associated with the purpose of pumping blood. It will work the way it is supposed to work, where the right way of working is partly constituted by fulfilling its function in normal conditions.

It turns out, then, that the e-functional picture constitutes itself in a fairly straightforward norm-identification machinery; here is how: first, what we need is to take a look at the relevant function plausibly served by the trait/artefact/practice in question. Once the function is identified, the question we need to ask ourselves is: how did the trait/artefact/practice plausibly fulfil its function at the moment of function acquisition? The answer to this question will give us normal functioning which, on the etiological account, corresponds to proper functioning. In turn, the relevant trait/artefact/practice will be functioning permissibly if and only if it is properly functioning; like this, we get the content of the norm we are after.

Furthermore, on this picture, we also get an easy way to identify the type of norm at stake: norms will be typed by the corresponding functions, which, in turn, are typed by the produced benefit. The norm ‘hearts ought to beat at this-or-that rate’ is a biological norm in virtue of the fact that it serves a biological function—pumping blood—which, in turn, produces a biological benefit.

One question that one may ask at this point regards normative force: isn’t it the case that, on functionalism, norms and values come too cheaply? It seems there could be a norm in a certain country that everyone must sleep exactly six hours a night. If there is a norm, there must be a function, and a value, by the lights of functionalism. Does the account have the resources to distinguish norms with “normative force” proper and values proper from others? (The big question in the offing, of course, is whether epistemic norms have normative force proper).¹¹

The answer, in short is: the functionalist picture alone stays neutral on the normative force of the norms at stake, as well as on whether the relevant functions instantiate goods *simpliciter*, i.e. for their own sake, or mere goods for one particular purpose or person, or attributive/predicative goods and so on. Particular types of values will deliver particular types of functions, which will deliver particular types of norms. Compatibly, these values/norms might be strictly domain specific rather than goods for their own sake/categorical norms. Ernie Sosa’s (2007) defends a similar value-theoretic picture, which might help the reader to see the landscape I have in mind: in Sosa’s view too, there are such things as critical domains, which are associated with goods or values for that domain. Any critical domain has at least one central value. This central value is fundamental to the domain in the sense that the central value is valuable for its own sake, *relative to the domain*. Crucially, central

¹¹ Many thanks to two anonymous referees for pressing me on this. See also Wedgwood (2018) for a useful taxonomy of types of norms by normative force. See also (Hazlett, 2013) for pessimism and (Simion, 2020) for optimism on this issue.

values organise the evaluations in that domain in the sense that all other domain specific values and norms will be derivative from the central ones: they can be explained fully in terms of the central values. I think of practices and their intrinsic functions along similar lines: practices are critical domains and their intrinsic functions are values for their own sake relative to the domain in question.

The example of the six hour sleep, for instance, I take it, is an example where what we have is a social value/function in the relevant society—since it is associated with a social norm. It might not be a value in any other sense, of course—very often social values, unfortunately, are not. Similarly, efficient killing could be a domain-specific value in the domain of crime—which, in turn, could generate domain-internal functions and norms. It's a completely different question: is the domain itself valuable *simpliciter* to begin with?—I take it that the domain of crime is not. It is also a completely different question whether particular social values are also good *simpliciter*, or morally good etc. and thus deliver categorical normative force. Functionalism stays neutral on this, as it should: these are substantive domain specific-issues, not pertaining to the general theory of how norms and values are associated.

That being said, the view does have the resources to accommodate the intuition that epistemic norms are categorical norms in a way in which crime norms are not: after all, it is plausible that the epistemic domain is valuable for our survival as a species, in a way in which, again, plausibly, the domain of crime is not. If species survival delivers categoricity—which is an etiological functionalist-friendly way to think about things -, then epistemic norms are categorical norms.¹²

2.2 Epistemic functionalism

In line with traits and artefacts, our practices too have etiological functions. They are the functions that contribute to the explanation of the continuous existence of the practices in question. Take medicine: it is plausible that generating health is the etiological function of the practice: that's why we continue engaging in it.

I will dub the totality of our epistemically significant endeavours 'the practice of inquiry.' Nothing hinges on this: it's a technical term meant to encompass our epistemic endeavours, broadly conceived, in the same way in which 'medicine' encompasses our health-generating scientific endeavours, broadly conceived.¹³ The view I favour takes knowledge to be the function of our epistemic practice of inquiry. Several people in the literature have given compelling arguments in support of the claim that knowledge is the goal of inquiry (Kelp Forthcoming, 2021, 2014; Millar, 2011), the function of (some of) our epistemic practices (Turri In Press, SimionSimion,

¹² See also (Simion, 2020) for a more sophisticated version of this value-theoretic picture.

¹³ In this, 'inquiry' as used here will be a broader term than the folk usage of 'inquiry,' in that it will not necessitate conscious wondering or question asking. Forming beliefs about one's surroundings while walking down the street, for instance, will be a move in inquiry. See also (Kelp, 2014, MS) for a similar technical usage.

2021), the fundamental epistemic value (e.g. Williamson, 2000). I will not rehearse these arguments here.¹⁴

If the knowledge function claim is right for the practice of inquiry, we have all we need to move on to identifying particular epistemic norms that we are interested in. Here is why: by stipulation, I take epistemically significant actions and states such as believing, judging, asserting, reasoning¹⁵ etc. to be moves in the practice of inquiry. Further on, it is plausible that moves in practices aim to make progress towards fulfilling the function of the practice; diabetological consults, brain surgeries etc. aim at generating health, in virtue of being moves in the practice of medicine; accelerating, breaking, stopping at the red light aim at safely getting you to your destination in a reasonable amount of time in virtue of being moves in the practice of driving, and so on. If that is the case, moves in the practice of inquiry—that is, all epistemically significant states and actions—will be governed by norms borne out by this central knowledge function of the practice.¹⁶

In what follows, I will put the functionalist machinery to work in the service of developing an integrated view of the epistemic normativity of reasoning.

¹⁴ Here I just assume the knowledge—rather than, e.g. truth—function and rely on extant defenses to do so. Elsewhere (2020, 2021) I argue extensively for a knowledge function of inquiry and against a truth function on etiological grounds. Space does not allow me to outline the argument here, but here is one important thing to mention: say that generating true belief is the main function of inquiry (e.g. Wedgwood, 2002). Would it then be reasonable to suppose that reasoning is governed by a corresponding truth norm? The short answer is ‘no,’ at least not on an etiological functionalist picture. Even if the function of reasoning is generating truth, proper functioning will still amount to reasoning from knowledge; after all, on etiological functionalism, proper functioning is normal functioning, and normal functioning is the way in which the trait functioned back when it acquired its function. Given that knowledge is easily available (indeed, the vast majority of our true beliefs are knowledgeable), very plausibly, at the moment of function acquisition, reasoning generated true belief by proceeding from knowledgeable premises. If so, knowledge is the condition for the normal functioning of reasoning. See (Simion, 2020) for more on this, and for a lengthy defence of a knowledge rather than truth function of inquiry.

¹⁵ For now, I simply assume that reasoning (in general) is a move in inquiry. See below for arguments as to why we should think of practical reasoning in these terms as well.

¹⁶ To be more precise, moves in practices aim (either directly or indirectly) to fulfil the function of the practice. The difference between direct and indirect aiming lies with achievability: for some moves in practices, the goal of the practice is only indirectly rather than directly reachable. Diabetological consults are moves in the practice of medicine, and they aim directly at fulfilling the goal of the practice of medicine: curing diseases. In turn, some moves aim at this final practice goal only indirectly, while aiming directly at intermediate goals: performing blood sugar tests aims (directly) at informing the diabetologist as to how well the patient is doing, which, in turn, aims directly at her diagnosing the patient correctly, and, further, at curing her disease. In this, blood sugar tests aim indirectly at the goal of the practice of medicine: curing diseases. They aim at making progress towards it. Similarly, since beliefs, assertions and conclusions of reasonings can be knowledgeable, in a way in which things like credences and withholdings cannot, I find it plausible that, e.g., belief formation aims directly at the aim of the practice (knowledge), while, at the same time, credence formation plausibly aims at it indirectly: on the face of it (although further research is called for on the issue), it seems to me that credence formation aims directly at adjusting ones doxastic states to the available evidence, which, in turn, plausibly ultimately aims at the aim of inquiry—knowledge generation.

3 No epistemic norm for action

3.1 Action, epistemic norms, and norms with epistemic content

One central debate in recent literature on epistemic normativity concerns the epistemic norm for action. Several people think that, where one's choice is p -dependent, it is appropriate to treat the proposition that p as a reason for acting only if one knows that p ¹⁷ (e.g. Hawthorne and Stanley (2008)). The most notable competing account puts forth a Bayesian norm, according to which it is rational to choose an act only if it maximizes expected utility (e.g. Douven, 2008).¹⁸

I have argued in previous work (Simion, 2018) that there is no epistemic norm for action: since norms of a particular type T uncontroversially bear an explanatory relation to values of type T, and since action in general is not thus associated with epistemic goods, I argue, epistemic norms do not govern action in general. Compatibly, epistemic norms can regulate specific actions that bear association relations to epistemic values, such as e.g. speech acts. Also compatibly, action in general will be governed by prudential norms with epistemic content, i.e. norms asking for more or less *epistemic* support for *prudentially* permissible action. I will not rehearse these arguments here. Rather, I will restrict myself to applying the normative framework outlined in the previous section to the normativity of action. We will see that, somewhat unsurprisingly, functionalist normativity will offer further support to this picture.

O the normative framework described in the previous section, a norm's pertaining to one type or another has to do with the function it is borne out by. Thus, prudential norms will be borne out by prudential functions—e.g. maximizing practical utility—, while epistemic norms will be concerned with reliably fulfilling epistemic functions in normal conditions.

If that is the case, however, it is not clear why action in general would be governed by an epistemic norm to begin with. When does it make sense, on this picture, to regulate something X by a norm of type Y, such as a prudential, moral or epistemic norm? When X has a function of type Y. Consider antibiotics; the function of producing them is curing bacterial infections. As such, norms governing this activity will plausibly be there to insure that they reliably do so. Conversely, when X does not have a type Y function, there is little reason to regulate X by a norm of type Y. Producing antibiotics will most likely not be governed by, say, aesthetic norms, given that the function of antibiotics is not to aesthetically please the consumer; antibiotics can be proper antibiotics even if they are not particularly pretty.

Thus, no epistemic category seems to meet E-Function for action: there is no epistemic good that action in general has generated in the past, and which contributes to the explanation of why we continue to engage in actions. Rather, prudential goods will meet E-Function for action: what explains our continuous engaging

¹⁷ While I will focus on the necessity direction of the norm, everything I say goes for the sufficiency direction, and therefore the biconditional version, too.

¹⁸ For an overview of the debate, see (Benton, 2014). For recent work, see (Littlejohn and Turri, 2014).

in action is its prudential value. Notice that it is also intuitive that most actions do not have epistemic functions; my eating breakfast, running in the park, brushing my teeth, buying chocolate, helping my old neighbour cross the street, are cases in point. Most of them are aimed at prudential goals, such as maximizing expected practical utility, some of them at moral goals, maybe a few at aesthetic goals. In the absence of any characteristic epistemic function, though, there is little reason to think that these actions will be governed by an epistemic norm.

Consider, in contrast, asserting, perceiving, reporting, judging, learning, reading, applying to the university etc. These actions all have epistemic functions. As such, it makes sense for them to be governed by epistemic norms. Having an epistemic norm for action makes perfect sense when it comes to actions that have epistemic functions, like assertion. It makes sense to ask what property exactly one's assertion must enjoy for it to be *epistemically* proper; that is, properly equipped to fulfil its epistemic function. However, just as in the case of producing antibiotics, there is no reason to think that my buying chocolate will be governed by an epistemic norm, due to the fact that it is not aimed at delivering epistemic goods.¹⁹

To sum up: epistemic norms regulate specific actions that bear association relations to epistemic values, such as e.g. speech acts. Compatibly, action in general is not governed by epistemic norms, but rather by prudential norms with epistemic content, i.e. norms asking for more or less epistemic support for *prudentially* permissible action. Just because a norm has epistemic content, it need not follow it is an epistemic norm (Simion, 2018). Consider: 'Only jump in the lake if you know how to swim!', in contrast with 'Assert that p only if you know that p!.' The latter is an epistemic norm with epistemic content: it tells you what epistemic support you need in order to *epistemically* permissibly assert. The former, in contrast, is merely a prudential norm with epistemic content: it tells you what epistemic support you need in order to *prudentially* permissibly jump in the lake. The norm for action, I contend, that the literature has been searching for, is a prudential norm with epistemic content, in virtue of bearing an association relation with a prudential rather than an epistemic function.

3.2 Action and practical reasoning

Note that nothing I have said so far concerns the epistemic norm for practical reasoning; while action is not governed by an epistemic norm, it may still be that practical reasoning is, insofar as it serves an epistemic function.

It is important to note that, even if practical reasoning does turn out to have an epistemic function, and therefore will be governed by an epistemic norm, this will not transmit epistemic requirements to the generated action. And here is why: I take it that

¹⁹ One interesting by-product of this result concerns the assertion-action commonality assumption; that is, the fairly popular claim that, given that assertion is a type of action, the epistemic norm governing the former is going to be but an instance of the epistemic norm governing the latter. Notice that the above results undermine commonality, together with the motivation behind it: assertion is not governed by an epistemic norm in virtue of its being a type of action, but due to its characteristic epistemic function.

the transmission thought would go along the following lines: since one ought to know that p in order to use p in practical reasoning, and since one ought to only act on the conclusions of permissible practical reasoning, it follows one ought to know that p in order to act. Now, the problem is that deontic transmission need not preserve the type of permissibility at stake in the premise of interest to us. To see this, consider, again: one ought to drive 30 mph in order to get safely at one's destination (traffic norm); one ought to get safely at one's destination in order to keep one's promise to mom (moral norm); Therefore, one ought to drive 30 mph in order to keep one's promise to mom (moral norm). Even if the first premise states a traffic norm, it need not follow that the conclusion does too. Similarly, in the argument above, even if the first premise states an epistemic norm, it does not follow the conclusion does too.

One other crucial thing to notice is that, if I am right about there not being an epistemic norm governing action in general, we need to substantively revise our methodology when we search for the epistemic norm of practical reasoning. Here is why: most cases put forth in defence or one account or another in this literature appeal to intuitions about propriety of acting in a particular situation. If this paper is right, however, the latter says little about the *epistemic* propriety of *practical reasoning*: prudential propriety of an instance of practical reasoning will likely depend on the prudential propriety of the generated piece of action; importantly, though, epistemic propriety need not. Thus, arguably, on pain of deontic equivocation, this methodology needs to be revisited: arguments pro and against one account or another of the epistemic norm for practical reasoning will have to stay clear of considerations pertaining to the prudential propriety of the corresponding action.

The next section does just that: it builds a defence of an epistemic norm of practical reasoning that relies on theoretical considerations about its function in the epistemic practice of inquiry, and thus stays clear of intuitions about action propriety.

4 The epistemic function and epistemic norm of reasoning

4.1 Knowledge and good reasoning

The view I want to defend when it comes to the epistemic function of reasoning (in general, be it theoretical or practical) is that it amounts to generating knowledge.²⁰ Here are some reasons in support of this thought:

First, since we are working with a view on which the practice of inquiry aims at generating knowledge, and moves in practices aim at making progress towards fulfilling the aim of the practice, and since reasoning is a move in the practice of inquiry, what my normative picture straightforwardly²¹ delivers is a knowledge-generating function of reasoning.

²⁰ See (McHugh and Way, 2018) for an excellent defense of an account of reasoning as a functional good-making kind.

²¹ This is a bit fast, of course; we have seen already that some moves in practices only aim indirectly at its goal. For our purposes here this will not make much difference.

Second, this view is also supported by the popular thought that good belief is knowledgeable belief. That's what it is for a belief to be a good token of its type: a belief that is not knowledgeable is defective *qua* belief. Beliefs that fall short of knowledge are criticizable (which, importantly, is not to say that the believers in question are: to the contrary, very plausibly, justified believers are not criticisable). In turn, attributive criticisability—i.e. criticizability of an X as a token of its type—implies evaluative norm violation: it cannot be that something X is criticisable in the attributive sense although there is nothing attributively defective with X. If all this is so, (5) follows. As Williamson puts it, mere belief is botched knowledge (2000). Crucially, note also that knowledge is enough for non-defective belief: there's nothing wrong with a knowledgeable belief that falls short of certainty, for instance. If so, on the assumption that the output of reasoning is a conclusion belief, and since beliefs ought to be knowledgeable, it follows that the function of reasoning is to generate knowledgeable conclusion beliefs.

Third, note that knowledge meets E-Function for reasoning: instances of reasoning have delivered knowledge in the past (empirically highly plausible), this benefitted us epistemically (the widely endorsed value of knowledge thesis),²² which contributes to the explanation of why we keep engaging in reasoning (empirically highly plausible). If this is right, we can straightforwardly employ our functionalist machinery to the aim of figuring out what epistemic norm is borne out by this picture.

What epistemic norm of reasoning will drop out of the knowledge function? Recall that in order to identify the corresponding norm, one needs to look back to when reasoning acquired its function, and ask the question: how did it work back then? How did it generate knowledge of its conclusion? I submit that the overwhelmingly plausible answer, given the easy availability of knowledge, and given that, in the vast majority of instances of reasoning, knowledge of the premises is needed for achieving knowledge of the conclusion, is that reasoning generated knowledge of the conclusion from known premises. If this is right, knowledge is the epistemic norm of reasoning:

KNR: Relying on p as a premise in reasoning is permissible iff one knows that p.

One question that arises at this point concerns the scope of the knowledge function and knowledge norm claims: is this result restricted to theoretical reasoning, or will it apply to practical reasoning as well?

The next section develops a unified view of the epistemic function and normativity of practical and theoretical reasoning. In order to do this, it outlines my preferred view of how practical and theoretical reasoning relate to each other.

4.2 Practical vs. theoretical reasoning

There are notable difficulties in offering a precise recipe for sharply distinguishing practical from theoretical reasoning. One fairly straightforward way to go about it

²² See e.g. Kelp and Simion (2017).

is to take theoretical reasoning to be concerned with answering descriptive questions (about matters of fact), while practical reasoning aims at answering distinctively normative questions (about what one ought to do).²³ The problem with this straightforward recipe is that it is too simple. Here is why: on one hand, one can think that there are such things as normative matters of fact, i.e. facts about what one ought to do. Conversely, it is plausible to think that conclusions of theoretical reasoning carry normative power too: after all, it looks as though, very roughly, proper instances of theoretical reasoning make it permissible to believe their conclusions, while improper instances thereof deliver impermissible conclusion-beliefs. It is tempting, then, to think that theoretical reasoning too attempts to answer a normative question: the question of what one should believe (Moran, 2001). Seen in this way, the contrast between practical and theoretical reason becomes essentially a contrast in target of normativity—i.e. belief vs. action: practical reason is concerned with answering the question ‘what should I do,’ while theoretical reasoning is concerned with answering the question: ‘what should I believe.’ Let’s call this way to distinguish practical from theoretical reasoning ‘the Simple View.’

Some philosophers are disenchanted with the Simple View: they think that there is more to the distinction between practical and theoretical reasoning than the Simple View suggests. In particular, they think that the two forms of reasoning differ not only in their target, but, importantly, also in their consequences: theoretical reasoning produces changes in one’s overall set of beliefs, whereas practical reasoning gives rise to intentional action; it is practical not only in its subject matter, but also in its issue. According to the defenders of this view, then, practical reasoning leads to modifications of our beliefs, whereas practical reasoning leads to modifications of our intentions ((Harman 1986), (Bratman 1987), henceforth also the Complex View).

I sympathise with the Simple View.²⁴ Here is why: as defenders of the Complex View themselves acknowledge, a more accurate way to represent the consequences of practical reasoning would be to say that deliberation about action generates appropriate intentions insofar as an agent is rational (Korsgaard, 1996). After all, many of us display weakness of will. But if that is the case, intention seems to be a contingent consequence of practical reasoning, premised on the agent’s rationality. In other words, it’s a consequence practical reasoning *should* have, rather than one that it essentially has, in virtue of the type of reasoning that it is. To see this, note that we don’t want to say that, in cases of weakness of will, no practical reasoning was exercised, in virtue of no intention being generated. To the contrary, what plausibly goes wrong in cases of weakness of will is precisely that an instance of practical reasoning is at work, but the agent fails to form the intention corresponding to its conclusion. In other words, the agent knows what she should do but fails to form the intention to do it.

²³ See (Wallace, 2020) for a nice overview.

²⁴ While I have strong sympathies for the Simple View, my account does not, strictly speaking, hang on it. If it turns out that (what I call below) the Revised Complex View is correct, my view will predict that (1) there is no epistemic norm for practical reasoning, since it doesn’t have any epistemic function, and (2) reasoning about what one ought to do is an instance of theoretical reasoning, governed by a knowledge norm.

If this is the case, a more plausible picture is the one outlined by the Simple View. Compatibly, defenders of the Complex View can come back with a weaker proposal, on which they supplement the Simple View with a normative rather than a descriptive claim: On the Revised Complex View, an instance of reasoning will be an instance of practical reasoning just in case its target of application is action, and it *should* deliver a change in intention. I'm not particularly convinced by this version of the Revised Complex View either; after all, we can easily imagine instances of practical reasoning which deliver the result that some intention that I already had was correct; in this case, no change in intention should follow. That being said, I don't need to settle this issue here, and I'm confident that more plausible ways to spell out the Complex View are available. The view I develop below is perfectly compatible with both the Simple and the Complex View. Indeed, I believe it offers a nice compromise between them.

According to the view I favour, there is no distinction between practical and theoretical reasoning when it comes to the nature of their conclusion-attitudes²⁵ (as per the Simple View): in both cases, the conclusion is a belief. The crucial distinction amounts to – in line with the Complex View—the *prudential* normative pressure that the conclusion exerts over one's intentions and actions: in the case of practical reason, one's actions and intentions (*prudentially*) should align with the conclusion belief: one (prudentially) should intend to do and do what one believes one ought to do as a result to one's good practical reasoning.

We have already seen that, on the picture defended here, good belief is knowledgeable belief. As such, on this picture, in contrast to action, practical reasoning does also serve an epistemic function, on top of its prudential function: generating a good belief—i.e. knowledge—of what one ought to do. Note, also, that knowledge meets E-Function: instances of practical reasoning have delivered knowledge of what one ought to do in the past (empirically highly plausible), this benefitted us epistemically (value of knowledge thesis), which contributes to the explanation of why we keep engaging in practical reasoning (the argument for this claim will parallel the one developed for knowledge as the function of inquiry in Chapter #2, so I will not rehearse it here). If that is so, practical reasoning will, arguably, afford an epistemic norm.

But is this really *a* function of practical reasoning, one might wonder? Isn't *the* function rather to generate an intention and/or the corresponding action, as the (Revised) Complex View would have it? And if this is so, will it not be the case that the arguments presented here against there being an epistemic norm for action apply, *mutatis mutandis*, to practical reasoning too?

A few things about this. First, note that the view defended here accepts that practical reasoning does have an important prudential function in generating prudentially proper intentions and actions. This prudential function will generate a prudential norm for practical reasoning, which may well have epistemic content. It is not my business here to discuss what the prudential norms for action or practical reasoning might be, I will leave that to practical philosophy to decide. What I am interested in is whether practical reasoning also has an epistemic function—and my answer is

²⁵ I follow (Broome, 2013) in using this terminology.

yes: to produce epistemically good conclusion beliefs—i.e. knowledgeable beliefs.²⁶ Importantly, the Complex View does not exclude the possibility of practical reasoning *also* generating a conclusion-belief on top of the relevant intention (indeed, the possibility of weakness of will requires there to be something else that plays the role of the conclusion of the relevant instance of practical reasoning; belief seems well qualified to do so). All they need for their view is that practical, but not theoretical reasoning, *also* has/should have intention as a consequence.

Second, I would worry a bit about a picture on which knowledge of the conclusion has nothing to do with the propriety (of sorts) of practical reasoning, and the latter is only to be assessed in terms of the propriety of the resulting intention. To see why, consider a case in which I reason as follows: “It’s raining outside/If it’s raining, I ought to take an umbrella/Therefore, I ought not take an umbrella”. This is a bad instance of reasoning. Now, say that, in spite of my conclusion belief, I go ahead and take the umbrella anyway. A picture on which the prudential function is the only function of practical reasoning—in virtue of its only consequence being the corresponding intention—has trouble explaining what went wrong here. Here is why: a view like this will have to describe the case above as one where two pieces of reasoning are happening; one theoretical (the one with the bad conclusion belief) and one practical (the one with the good conclusion-intention). After all, this view takes the output of the relevant piece of reasoning to determine the nature thereof—theoretical or practical. So what we have above on this view is a correct instance of practical reasoning (from it’s raining and if it’s raining, I ought to take the umbrella to actually forming the intention of taking it and, in fact, taking it) and a bad instance of theoretical reasoning (to the conclusion belief that I ought not take the umbrella). This picture, however, fails to explain the following datum: it looks as though, my taking the umbrella after I reach the (mistaken) conclusion that I ought not do it makes things worse rather than better. After all, not only am I a bad reasoner, but my actions and my beliefs also seem to come in conflict with each other.

But couldn’t the defender of this view also explain away the intuition of impermissibility in this example by appeal to requirements of structural rationality (such as, e.g., akrasia), independently of without also assuming further epistemic functions for practical reasoning? Furthermore, it would seem that the datum could also be explained by the negative prudential consequences that would ensue from holding conflicting attitudes of this sort.²⁷

The answer, in short, is ‘no’: the problem doesn’t go away if we stipulate that there is no intention to take the umbrella formed in this case, but rather just the corresponding action—thus no attitudes to come in conflict.

Also, even if it is plausibly true that the prudential function is *the main* function of practical reasoning, it is perfectly compatible with practical reasoning having a

²⁶ (Dancy, 2018) also takes practical reasoning to be governed by both practical and epistemic norms, and takes the practical to take primacy. While it falls beyond the scope of this paper to discuss the priority issue, I want to flag that I do agree with Dancy on this front. See also (Simion, 2020) for a book-length defence of a functionalist normative picture that bears this out.

²⁷ Many thanks to an anonymous referee for pressing me on this.

prudential function (generating a prudentially good intention/intentional action) that it *also* has an epistemic function (generating knowledge of what one ought to do). Insofar as an epistemic function is present, an epistemic norm will drop out of it.

If I am right in thinking that theoretical and practical reasoning serve the same epistemic function – generating knowledge of the conclusion – the argument for the norm governing reasoning presented in the previous section will apply to practical reasoning as well. What we get, then, is a knowledge norm for practical reasoning:

K-Norm of Practical Reasoning: Relying on p as a premise in practical reasoning is permissible iff one knows that p .

Note that, importantly, the functionalist claim does not imply that knowledge of premises is always necessary for generating knowledge of the conclusion. After all, functionalist norms do not track necessary conditions for function fulfilment, but merely the most reliable ways to fulfil the corresponding functions in normal conditions. The claim is that reasoning from premises one does not know is always impermissible, although it might, on occasion, lead to function fulfilment.

To see why this an important advantage of the view, consider the following variation of the umbrella case above: The probability of rain is 0.98/If there is a 0.98 probability of rain, I should take an umbrella/Therefore, I should take an umbrella. Now, say that I'm wrong: the probability of rain is actually 0.97. Still, intuitively, I know the conclusion: I should take an umbrella.

Consider, also, a corresponding case of knowledge from falsehood²⁸ in the case of theoretical reasoning: Counting with some care the number of people present at my talk, I reason: 'There are 53 people at my talk; therefore my 100 hand-out copies are sufficient'. My premise is false. There are 52 people in attendance—I double counted one person who changed seats during the count. And yet, intuitively, I know my conclusion.

Cases like this are thought by many²⁹ to be very problematic for knowledge norms of practical and, respectively theoretical reasoning. Crucially, they are not problematic for the view defended here: just like hearts can pump blood even when improperly functioning, instances of reasoning can deliver knowledge when improper too. This, however, does not falsify the normative claim that reasoning should always proceed from knowledge.

In further good news, the picture defended here is also compatible with cases of intuitively permissible action (and intention) from (epistemically) bad reasoning. We are fallible creatures: We often rely on heuristics and habits when we act, and, short of disaster, that's perfectly fine: we have limited computing capacities. Also, we may often not even consider the option that corresponds to what we ought to do. We often also get our priorities wrong. We do things in the wrong order. However, we still get everything done (although perhaps with more effort) or at least well enough to avoid disabling disaster. Again this may happen quite often. On pain of extreme demandingness, one could argue, knowledge of oughts does not seem needed for permissible action in all walks of life. I don't want or need to take a stance on this

²⁸ See e.g. (Luzzi, 2019) for an excellent discussion.

²⁹ See (Weisberg, 2013).

issue: my view is perfectly compatible with the prudential norm governing practical reasoning (and action) being much weaker than knowledge, or even stakes-variant. All my view requires is that knowledge of the premises be present for the *epistemic* norm of practical reasoning to be respected, which, in turn, in normal conditions, will reliably generate epistemic function fulfilment: knowledge of the conclusion. Compatibly, the epistemic norm may well be often overridden by prudential considerations, in which case less than knowledge will be needed for prudentially permissible practical reasoning.³⁰ That's fine (all-things-considered).

Here is one last worry I would like to address: Why should we assume that practical reasoning also has an epistemic function, in addition to its practical (/moral etc.) one to begin with? Can't we explain everything we need to explain by assuming that practical reasoning is governed by a prudential knowledge norm?³¹ Couldn't its prudential function alone explain why one is permitted to use *p* as a premise in practical reasoning if and only if one knows *p*? For instance, couldn't it be the case that reasoning and forming intentions only on the basis of knowledge would guarantee better practical results in the long run and safer achievements (Fassio, 2017)?

Importantly, although I don't defend this view myself, my view is compatible with a view according to which the prudential etc. function of action and practical reasoning also predicts that knowledge is the prudential norm for action and practical reasoning, sourced in their prudential function. If this turns out to be the case, on my view, the epistemic and prudential norm will just happen to have the same content, in virtue of the same condition being reliably conducive to both the relevant prudential etc. goods/function fulfilment and the relevant epistemic goods/function fulfilment. However, my view also has the flexibility to allow that this is not the case, and thus to explain both cases put forth in the literature in support of a knowledge norm, and cases put forth to show that knowledge is too strong a norm for some cases (e.g. umbrella cases) and too weak for others (e.g. high stakes cases):

³⁰ A referee worries about conflicts of norms on my view. The view commits me to the following claims: In high stakes cases, it is epistemically permissible for me to reason from *p* insofar as I know that *p*. This might result in my affirming conclusion propositions like *the best option is to do X*. But prudentially I ought not reason from this. The referee worries that this is an odd result because these cases don't feel like the standard sort of cases in which one has good evidence for *p* but has been threatened not to believe *p*.

I agree with the referee that the cases are not clearly analogous: to see what is going on, we will need to distinguish between two high stakes cases: In Case 1, my practical reasoning will not lead to action: I engage in it for the fun of it. Case 2 is a normal, everyday case of practical reasoning, whereby the probability of my acting on the conclusion belief is quite high. On my view, in the first case, relying on *p* is both epistemically and prudentially permissible: I will come to know the conclusion if I do so, and nothing bad is likely to happen to me as a result. In Case 2, however, the risk of my acting without enough epistemic support given the high stakes makes the corresponding instance of practical reasoning prudentially impermissible, although epistemically perfectly fine. I think this is the right result; to see this, consider a variation on a famous case by Jessica Brown (2010): If I know that my believing that my boss is bald will trigger my compulsive assertion to his face that he is bald, I'd better abstain from reasoning myself to this conclusion. That's not because I would be thereby breaking epistemic norms, but rather because I want to keep my job. Thanks to the referee for pressing me on this.

³¹ Many thanks to an anonymous referee for pressing me on this.

it can do so by allowing variability in the prudential norm, while, at the same time, keeping the content of the epistemic norm fixed.

4.3 Conclusion

The debate concerning the epistemic status a premise needs to enjoy in order to be permissibly employed in practical reasoning features two main actors—the knowledge norm defenders for both action and practical reasoning on one hand (e.g. Hawthorne & Stanley, 2008; Williamson, 2000), and the stakes-variant norm defenders for both action and practical reasoning on the other (e.g. Brown, 2010, 2012; Douven, 2008).

This paper proposed a picture on which action and practical reasoning are, to some extent normative strangers: they share a prudential function, but not an epistemic function; as such, they are both governed by prudential norms, but practical reasoning, and not action, is also governed by an epistemic norm in virtue of its epistemic function of generating knowledge of what one ought to do. Excitingly, though, this picture unifies the epistemic normativity of reasoning: practical and theoretical reasoning share epistemic norm in virtue of sharing epistemic function: generating knowledge of the conclusion.

Acknowledgements Many thanks to Jessica Brown, Davide Fassio, Mikkel Gerken, Sandy Goldberg, Peter Graham, Patrick Greenough, Chris Kelp, Matt McGrath, Conor McHugh, Andrew Reisner, Jonathan Way, Daniel Whiting, Asa Wikforss, Tim Williamson, and two anonymous referees for Synthese for their help with this paper. This research was funded by the Mind Association (Mind Fellowship for the ‘Epistemic Norms and Epistemic Functions’ project, 2018–2019) and the European Research Council (KnowledgeLab: Knowledge First Social Epistemology project, grant agreement 948,356).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Benton, M. (2014). Knowledge Norms. *Internet Encyclopedia of Philosophy*.
- Bratman, M. (1987). *Intention, Plans, and Practical Reason*. Cambridge, Mass.: Harvard University Press.
- Broome, J. (2013). *Rationality through reasoning*. Wiley-Blackwell.
- Brown, J. (2010). Knowledge and Assertion. *Philosophy and Phenomenological Research*, 81, 549–566.
- Brown, J. (2012). Assertion and practical reasoning: Common or divergent epistemic standards? *Philosophy and Phenomenological Research*, 84(1), 123–157.
- Buller, D. J. (1998). Etiological theories of function: A geographical survey. *Biology and Philosophy*, 13(4), 505–527.
- Cummins, R. (1975). Functional Analysis. *The Journal of Philosophy*, 72(20), 741–765.
- Dancy, J. (2004). *Ethics without principles*. Oxford University Press.
- Dancy, J. (2018). *Practical shape: A theory of practical reasoning*. Oxford: Oxford University Press.

- Douven, I. (2008). Knowledge and practical reasoning. *Dialectica*, 62(1), 101–118.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford University Press.
- Fantl, J., & McGrath, M. (2012). Pragmatic encroachment: It's not just about knowledge. *Episteme*, 9, 27–42.
- Fassio, D. (2017). Is there an epistemic norm of practical reasoning? *Philosophical Studies*, 174, 2137–2166.
- Graham, P. J. (2012). Epistemic entitlement. *Nous*, 46(3), 449–482.
- Graham, P. J. (2014). Functions. In A. Fairweather & O. Flanagan (Eds.), *warrant, history* (pp. 15–35). Virtue. Cambridge University Press.
- Harman, G. (1986). *Change in View*. Cambridge, Mass.: MIT Press.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105, 571–590.
- Hazlett, A. (2013). *A luxury of the understanding: On the value of true belief*. Oxford University Press.
- Kelp, C. (2014). Two for the knowledge goal of inquiry. *American Philosophical Quarterly*, 51, 227–232.
- Kelp, C. (2021). *Inquiry, knowledge, and understanding*. Oxford University Press.
- Kelp, C. & Simion, M. (Forthcoming). *Sharing Knowledge: A Functional Account of Assertion*. Cambridge: Cambridge University Press.
- Kelp, C., & Simion, M. (2017). Commodious knowledge. *Synthese*, 194(5), 1487–1502.
- Korsgaard, C. (1996). 'Skepticism about practical reason', in *her creating the kingdom of ends*. Cambridge: Cambridge University Press.
- Littlejohn, C. (2014). The unity of reason. In C. Littlejohn & J. Turri (Eds.), *Epistemic norms: Assertion, action, and belief*. Oxford: Oxford University Press.
- Luzzi, F. (2019). *Knowledge from non-knowledge: Inference, testimony, and memory*. Cambridge University Press.
- McHugh, C., & Way, J. (2008). What is Good Reasoning? *Philosophy and Phenomenological Research*, 96, 153–174.
- McHugh, C., & Way, J. (2018). What is Reasoning? *Mind*, 127(505), 167–196.
- Millar, A. (2011). Why knowledge matters. *Aristotelian Society Supplementary*, 85(1), 63–81.
- Millikan, R. (1984). *Language, thought, and other biological categories*. The MIT Press.
- Moran, R. (2001). *Authority and Estrangement*. Princeton: Princeton University Press.
- Neander, K. (1991). The teleological notion of function. *Australasian Journal of Philosophy*, 69, 454–468.
- Neta, R. (2009). Treating something as a reason for action. *Noûs*, 43, 684–699.
- Reed, B. (2012). Resisting encroachment. *Philosophy and Phenomenological Research*, 85, 465–472.
- Simion, M. (2018). No epistemic norm for action. *American Philosophical Quarterly*, 55(3), 231–238.
- Simion, M. (2019). Knowledge-First Functionalism. In: Lisa Miracchi (ed.) *Philosophical Issues*. Online First.
- Simion, M. (2021). *Shifty speech and independent thought: Epistemic normativity in context*. Oxford University Press.
- Simion (2020). *Epistemic Norms and Epistemic Functions*. Book Manuscript.
- Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. 1). Oxford University Press.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford University Press.
- Sullivan-Bissett, E. (2017). Biological function and epistemic normativity. *Philosophical Explorations*, 20(1), 94–110.
- Wallace, R. J. (2020). Practical Reason. *The Stanford Encyclopedia of Philosophy* (Spring 2020 Edition), Edward N. Zalta (ed.), URL=<<https://plato.stanford.edu/archives/spr2020/entries/practical-reason/>>.
- Weatherson, B. (2012). Knowledge, bets, and interests. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions*. Oxford University Press.
- Wedgwood, R. (2002). The aim of belief. *Philosophical Perspectives*, 16, 267–297.
- Wedgwood, R. (2018). *The unity of normativity*. *Oxford handbook of reasons and normativity*. Oxford University Press.
- Weisberg, J. (2013). "Knowledge in Action." *Philosophers' Imprint* 13, 1–23.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.