Synthese (2021) 199:4673–4693
https://doi.org/10.1007/s11229-020-02996-3

THE EPISTEMIC SIGNIFICANCE OF NON-EPISTEMIC FACTORS

Epistemic evaluation and the need for 'impure' epistemic standards

Nikola Anna Kompa[1]

Received: 2 May 2019 / Accepted: 7 December 2020 / Published online: 23 January 2021
© The Author(s) 2021

Abstract
That knowledge ascriptions exhibit some form of sensitivity to context is uncontroversial. How best to account for the context-sensitivity at issue, however, is the topic of heated debates. A certain version of nonindexical contextualism seems to be a promising option. Even so, it is incumbent upon any contextualist account to explain in what way and to what extent the epistemic standard operative in a particular context of epistemic evaluation is affected by non-epistemic factors (such as practical interests). In this paper, I investigate how non-epistemic factors come into play when knowledge is ascribed. I argue that knowledge ascriptions often serve the purpose of providing actionable information. This, in turn, requires that epistemic interests be balanced against non-epistemic interests. Moreover, it raises the question of whose interests matter, those of the ascriber, the addressee (of the knowledge ascription), or the subject of ascription. Eventually, an answer to the question is suggested.

Keywords Nonindexical contextualism · Epistemic standards · Practical interests · Functions of knowledge ascriptions · Threshold problem, whose stakes? problem

> *It is essential to realize that 'true' and 'false', like 'free' and 'unfree', do not stand for anything simple at all; but for a general dimension of being a right or proper thing to say as opposed to a wrong thing, in these circumstances, to this audience, for these purposes and with these intentions.*
>
> *John Austin*

1 Introduction

That knowledge ascriptions exhibit some form of sensitivity to context or other is not controversial. What is controversial, however, is how best to account for the context-sensitivity at issue. I take nonindexical contextualism to be a promising route. In what follows, I will briefly sketch the basic idea of the version of nonindexical

✉ Nikola Anna Kompa
nkompa@uos.de

[1] Institute of Philosophy, Osnabrück University, Albrechtstrasse, 28a, 49076 Osnabrück, Germany

contextualism that I favour. Yet the aim is not to defend this version of contextualism but rather to develop it further. A question any contextualist has to answer is how and to what extent the epistemic standard operative in a particular context of epistemic evaluation is affected by non-epistemic factors (such as practical interests). I will, therefore, investigate more closely how epistemic standards are thus affected. Eventually, I will claim that non-epistemic factors.

(i) come into play when epistemic agents decide which (impure) epistemic aim to pursue in a particular context.
(ii) More specifically, to the extent that the function of knowledge ascriptions is to provide *actionable* information, epistemic interests and needs have to be balanced against non-epistemic interests and needs.
(iii) Yet two cases need to be distinguished, according to *whom* the information is supposed to be actionable for: (a) the subject of the ascription (the putative knower), or (b) the addressee of the ascription. That will help fix the epistemic standard operative in the context of ascription and solve what Stephen Grimm calls *the whose stakes? problem* (Grimm 2015).

2 The context sensitivity of knowledge ascriptions

Over the last couple of decades, contextualists of various brands came up with a set of vignettes to illustrate the claim that epistemic standards (—standards for how much it takes to count as a knower) may vary from one context to the next, while the variation is due to non-epistemic or practical factors. Keith DeRose's by now classic bank case (DeRose 1992) or Jason Stanley's gender-sensitive version of it commonly serve as examples (Stanley 2005; cf. also, e.g., Annis 1978 or Cohen 1999). —In light of these cases, the contextualist would like to diagnose a certain aim-and-purpose dependency in our practice of ascribing knowledge. Knowledge, she will say, is something that we attribute to one another in light of certain goals, needs, and demands. The epistemic standard—the standard someone has to live up to in order to count as a knower (on the assumption that p is true) in the context at hand—will be set differently in different contexts, depending on the goal pursued, the needs to be met, and the function the knowledge ascription is supposed to fulfill in achieving that goal or meeting those needs.[1] Moreover, ordinary speakers seem to (implicitly at least) acknowledge varying contextual requirements governing the use of 'know'. If they didn't, the contextualist vignettes should not have attracted any attention in the first place. That they did is evidence that speakers are sensitive to context-dependent epistemic standards.

How do we explain the intuition of shiftiness or context-sensitivity regarding our practice of ascribing knowledge? Some try to *explain it away* by means of an error theory that invokes either certain pragmatic mechanisms or insights from

[1] Talk of standards is a workaround; a fuller account would have to be more explicit about the various parameters of context (cf. Baumann 2016) and the ways in which different interests can be balanced against one another. To assume that this will easily result in a single standard suggests a precision (und uniqueness) that is, presumably, not to be had.

cognitive psychology. Speakers come out as being systematically mistaken about the truth conditions of knowledge ascriptions or as being subject to certain cognitive biases. More specifically, it has been suggested to model the context-sensitivity *pragmatically*, e.g., via certain types of warranted-assertibility-maneuvers according to which speakers systematically conflate truth- and assertibility-conditions (cf., e.g., Davis 2004, 2005, 2007, 2017; Brown 2006; Stei 2014; or Rysiew 2007, 2017; Kvart 2018; the idea goes back to DeRose 1999). More recently, attempts have been made to explain away the intuition of context-sensitivity or 'shiftiness' by drawing on work in *cognitive psychology* (cf. Nagel and Smith 2017 for an overview): the shiftiness-intuition is said to be due to a form of availability bias (Williamson 2005), focal bias (Gerken 2012, 2013), or increased epistemic anxiety (Nagel 2010; cf. also Nagel 2008), to name just a few popular accounts.

Alternatively, one might try to *semantically* model the context-sensitivity at issue (i) as a form of indexicality, the resulting position being a form of indexical contextualism (cf., e.g., Cohen 1999; DeRose 2000; Blome-Tillmann 2014); (ii) as a form of dependency on circumstances of evaluation, resulting in nonindexical contextualism (also known as moderate relativism; cf. MacFarlane 2005, 2007, 2009; Kompa 2002, 2005, 2014, 2015, 2016; Kölbel 2008; Récanati 2008); (iii) as a form of assessment sensitivity, issuing in assessment-sensitive relativism (cf. MacFarlane 2005, 2014); or by (iv) incorporating it into an Austinian semantic framework (Lawlor 2013).[2]

3 Nonindexical contextualism

I opt for what is commonly classified as a semantic account.[3] More specifically, I take nonindexical contextualism to be a promising route.[4] It avails itself of the Kaplanian distinction between contexts of utterance and circumstances of

[2] The Austinian semantics and the non-indexical semantics may be almost notational variants of one another.

[3] The account is semantic in that it concerns truth- or correctness-conditions. Yet it is utterances' truth- (or correctness-) conditions, not the truth-conditions of sentences in isolation that are at issue here. And an utterance's content may be the result of pragmatic completion or enrichment (Carston 2002; Récanati 2004). Also, on the truth-conditional account that I favor, there is no other, purely semantic, truth-evaluable *content* that is expressed independently of the utterance's content (Kompa 2012, 2014). Rather, semantics constrains interpretation, by providing something like a propositional radical (Bach 2006), or by simply providing lexical information that can then feed into the process of interpretation (Kompa 2017). Thus, given that it is an account of utterance truth (correctness), it may alternatively be classified as pragmatic.

[4] Here are some of the advantages nonindexical contextualism enjoys over indexical contextualism. Nonindexical contextualism helps us to get rid of the problem of lost disagreement: speakers no longer talk past each other as the one affirms what the other denies (Kompa 2015). Also, nonindexical contextualists do not have to attribute semantic blindness (Davis 2007) to seemingly competent speakers. And they can account for retraction as a speaker might say "My previous utterance was correct or appropriate (back then), but X (really) doesn't known". Furthermore, the account provides a plausible semantics for evaluative terms in general; when speaker employ evaluative terms, they may disagree over meaning or over standards; similarly so in the case of "know", in which speakers might disagree over *what* it takes to count as a knower or over *how much* it takes to thus count (Kompa 2016).

evaluation (Kaplan 1989). Given the Kaplanian (or Lewisean) framework, context can affect an utterance's truth-value either by affecting which content has been expressed or by providing circumstances of evaluation on which the utterance's truth-value may also depend. Consequently, variation in an utterance's truth-value can either be due to a form of indexicality that results in variation in the utterance's content; or it can be due to a form of circumstance-sensitivity as the content gets evaluated differently in light of different circumstances of evaluation (or at different indices, Lewis 1980), which are in turn determined by the context of utterance. Which option the contextualist chooses determines whether she is an indexical or a nonindexical contextualist.

According to the latter, then, a context of utterance has two duties to perform. It has to supply the denotations to indexical and demonstrative terms in a sentence and to determine the denotation of otherwise under-determinate terms (if such there be)—yet the verb 'know' is not counted among them (—that is what distinguishes nonindexical from indexical contextualists). It thus helps to determine the utterance's content, i.e. the proposition expressed. It also has to provide features on which the truth of the whole utterance may depend, such as, arguably, an epistemic standard. The following necessary condition on utterance truth results:

> NIC: "An utterance of a sentence s (of the form 'A knows that P') in context c is true only if the proposition thereby expressed is true relative to circumstances (w_c, e_c), where w_c is the world of c and e_c is the epistemic standard operative in c" (Kompa 2015, p. 143).

The condition employs two distinct notions of truth: an un-relativized notion of utterance truth, and a relativized notion of propositional truth. In order to indicate the difference and to make clear that what is suggested is a type of utterance semantics, I will henceforth speak of utterances in general, and of knowledge ascriptions in particular, as being *correct* instead. A knowledge ascription is correct, then, only if the proposition expressed is true as evaluated relative to the epistemic standard operative in the context of ascription. For what other standard should be relevant to the *correctness* of a given knowledge ascription than the one operative in the context of ascription?[5] —This is the semantic story (yet cf. fn. 3); it concerns correctness conditions of knowledge ascriptions understood as utterances (or speech acts). But it raises the questions of (i) what standard is operative in a context of ascription, and (ii) how the standard is to be determined (i.e., *the whose stakes? problem*). I will address the two questions in turn. This is also where non-epistemic factors come into play.

[5] If its correctness depended on a context of assessment, as those defending assessment sensitivity claim (MacFarlane 2014), then how could a speaker even intend to make a correct utterance, give that there will be contexts of assessment relative to which it will come out true and others relative to which it will come out false (cf. Kompa 2015)?

4 The threshold problem

Concerning the first question, one might wonder whether there isn't a single, objectively correct and invariant epistemic standard that has to be met for a knowledge ascription to be true. Invariantists of skeptical leanings take the skeptical standard to be the only legitimate standard; it requires truth-guaranteeing justification. The concept of knowledge is reserved for that kind of ascription (or rather denial), they say; and our semantic theory should reflect that fact. As has often been pointed out, that would make all everyday knowledge ascriptions and even all ascriptions of scientific knowledge come out false. I side with those who take this to be a *reductio* of the skeptical invariantist semantics. Most invariantists agree anyway and aim at a more moderate standard. Yet pending an account of where exactly to set the invariant standard, this looks like a rather dogmatic (or desperate) move—a problem that goes by the name 'threshold problem' (cf., e.g., Hannon 2017). The question lurking in the background here is not *what* it takes to count as a knower (i.e., which epistemic property has to be instantiated) but *how much* it takes to thus count (Grimm 2015; Kompa 2016).

Some progress has recently been made. Krista Lawlor's reasonable person's standard, for example, is a promising idea (Lawlor 2013). But it is not an invariant standard, as it varies with the interests of the speaker and the addressee: the speaker, in making a knowledge ascription (and in thereby giving assurance), ought to take the addressee's interests into account (cf. Lawlor 2013, p. 160). It fits in nicely with David Henderson's suggestion that people's epistemic needs inform our judgments about who knows what (cf. Henderson and Horgan 2015, p. 95) and that epistemic norms qua social norms reflect (or guarantee) the agents' sensitivity to each other's epistemic needs (cf. Henderson 2020). It is an idea that any contextualist should be sympathetic to, and it will be taken up below. Note, however, that there are duties on both sides. The speaker has to take the addressee's (or the subject's, see the following) interests into account. Yet in case the addressee has 'special epistemic needs', it is incumbent upon her to make them manifest. If A asks B what time the opera begins and B says: "I know it begins at 8 o'clock", and then later finds out that A's life depends on it, he may well get back to B and say: "Please, don't take it from me" or "You better double-check; actually, I don't really know".

Proponents of Communal Impurism such as Michael Hannon (2017) or of Craigian Contextualism such as Robin McKenna (2013), on the other hand, hold that there is an *invariant*, communally determined threshold or standard for knowledge. Individual thresholds, although they can never drop below, can override the communal threshold; if much is at stake, for example. The tricky question is how and where the communal standard is to be set, then. We can all agree that it would be nice to have a lower epistemic bound or baseline–if only as a possible rejoinder to the cheap knowledge objection, according to which a community with very few epistemic needs and accordingly low epistemic standards would get knowledge almost for free. But that is not quite the standard at issue. Rather, the communally determined standard is conceived of as a standard "such

that many people can freely draw on" the information thus certified as knowledge (Hannon 2017, p. 614); it should "satisfy all comers" (Lawlor 2013, p. 3), or at least all reasonable ones. It will therefore be pretty demanding already and has to be "set fairly high" (Hannon 2017, p. 614).

Ordinary contextualists (like me), therefore, suggest that we resort to contextually varying standards. There is, according to the contextualist, no single, context-invariant epistemic standard that has to be met for a knowledge ascription to be true. No standard is singled out as the one-and-only epistemic standard. That is not to claim that standards vary dramatically and incessantly. In very common situations (in which people ask other people what time it is, where the station is, and similar mundane things), which are not 'epistemically marked' in any way, default standards may become effective. For all that, there are different yet equally legitimate 'impure' epistemic standards because there are different yet equally legitimate practico-epistemic goals one might pursue in ascribing knowledge. More fully, I would like to argue as follows:

1. An epistemic standard either accommodates practical interest or does not.
2. If an epistemic standard does not accommodate practical interest, the result will be a purely epistemic (and invariant) standard.
3. If an epistemic standard accommodates practical interest, and does so in a context-sensitive manner, the result will be a context-sensitive, 'impure', practico-epistemic standard
4. Knowledge ascriptions predominately serve the function of certifying actionable information.
5. Certifying actionable information requires that the respective epistemic standards accommodate practical interests in a context-sensitive manner.
6. Therefore, (felicitous) knowledge ascriptions predominately invoke context-sensitive, 'impure' practico-epistemic standards.

I take premise 1 and 3 to be unproblematic; premise 1 is stating the obvious, and premise 3 amounts to nothing more than a definition of "context-sensitive, impure practico-epistemic standard". In the following Sect. 5, I will discuss premise 2. Section 6 and 7 address premises 5 and 4 respectively. In Sect. 8, I will reply to some objections. Section 9, finally, suggests a solution to *the whose stakes? Problem* and thereby also examines the context-sensitive at issue in more detail.

5 Epistemic standards and epistemic goals

> Believe truth! Shun error! … these, we see, are two materially different laws; and by choosing between them we may end by coloring differently our whole intellectual life.
>
> William James

Knowledge ascriptions are epistemic evaluations; but evaluation seems to almost always take place against the backdrop of certain practical concerns and interests. If we abstract away from all practical interests, if we refrain from taking any practical interests into account, if we try to achieve the best possible epistemic state, we will arrive at a purely epistemic—and invariant—standard [—the case of the Higgs bosom discussed in Henderson (2020) is illuminating in this respect]. For purposes of illustration, imagine a purely contemplating creature with no need (nor interest or opportunity) to act in any way imaginable and with no practical limitations, a creature whose only interest was to improve epistemically. It has thus neither reason nor motivation to rest content with anything less than the finest epistemic position possible and to meet anything other than the purest epistemic standard possible.[6] Yet what would such a standard look like, that results from pursuing a purely epistemic goal? That depends, arguably, on which epistemic aim or goal one has. It is common to distinguish two different epistemic aims: maximizing true beliefs and minimizing false beliefs. Of course, there could also be a third (overarching) purely epistemic aim (understanding might be a promising candidate, cf., e.g., Kvanvig 2009a). Although I take this to be an interesting idea, I will leave it aside for the moment and stick to the more traditional aims. Now one might argue that said creature with no need to act (as it does not have any practical interests), and a fortiori no need to act before ALL the information is in (while also avoiding all possible errors), could pursue both aims at the same time, arriving at a purely epistemic standard, i.e., the standard that guarantees the best possible epistemic outcome (I will come back to this in Sect. 8, however).

Unfortunately, this is not the situation we (normal epistemic agents) commonly find ourselves in. We cannot pursue both aims with equal verve. We cannot wait until all the information is in while all error has been avoided. We need information, so that we can act upon it. Thus, when we maximize true beliefs, we already (have to) trade in quality for quantity, because we are, at bottom, practical (not purely contemplating) creatures; i.e. creatures that need to act and whose resources are limited. It has to be conceded that the two aims or goals impose conflicting requirements on us and cannot both be fully satisfied at the same time. As Wayne Riggs puts it:

> "Unfortunately, pursuit of these two goals can pull us in opposite directions. Concern to avoid having false beliefs naturally prompts skepticism and caution, while the desire to accumulate true beliefs urges us towards acceptance […]" (Riggs 2003, p. 343).

We have to choose which to value higher (in a particular context). And we will commonly do so on the basis of non-epistemic considerations, or so I will argue. On the assumption that there is no further, overarching epistemic goal, one could even claim that we HAVE to decide how to weigh them on non-epistemic grounds. However, I am not prepared to argue for the stronger claim here. All I am arguing is that these two goals are important epistemic goals and that non-epistemic factors such as practical interests commonly come into play in deciding which to value higher.

[6] Admittedly, and as Stephen Grimm notes, "the notion of what matters from a 'purely epistemic point of view' (or the like) is considerably less clear than the notion of what matters from a practical point of view" (Grimm 2015, p. 122).

6 Actionable information

So far, I have been claiming that practical interests and concerns come into play when epistemic agents decide which practico-epistemic aim to pursue in a given situation, as knowledge ascriptions are commonly undertaken with that aim in mind. This is not to deny that knowledge ascriptions serve epistemic aims; but they commonly also serve practical aims or interests (such as successful action). They are supposed to provide the addressee with epistemic, but also with practical reasons, i.e. reason to act (Robitzsch 2019).

By way of illustration, imagine two oracles.[7] The deliverances of the first are absolutely error-proof; all possible challenges have been met; it provides truth-guaranteeing justification. The deliverances of the second are less certain; various error possibilities aren't ruled out yet; it provides good but fallible justification. Which oracle would you prefer to consult—other things being equal? From an epistemic point of view, the choice seems to be straightforward.[8] The only reason you might have for choosing the second over the first is that the first will presumably have very little to say, while the second will deliver much more. But why more; why trade in quality for quantity? Because, often, a fairly good or robust epistemic position suffices. What does it suffice for? For all practical purposes! More often than not, what we need is actionable information. Of course, from a purely epistemic point of view, "[i]ncreasingly high levels of justification are valuable because they improve our cognitive situation…" (Hannon 2017, p. 608). Yet increasing effort is costly. Consequently, we commonly are epistemic satisficers. We weigh epistemic and non-epistemic needs and demands against each other, the result being an impure, practico-epistemic standard.

Mostly, people don't seem to pursue these overall epistemic aims (of maximizing true or minimizing false beliefs) anyway, but rather more specific, local epistemic aims, such as the aim of avoiding falsely negative or falsely positive judgments with respect to a particular subject matter (this somewhat relates to the debate on values in science; cf., e.g., Douglas 2000).[9] And it seems highly plausible to assume that these more specific epistemic aims (and the resulting standard) are also chosen on the basis of practical considerations, and with the intention of providing actionable information. Here are two cases to illustrate the point.

Climate Change Nowadays, it is often asked whether some particular person, group of persons or institution (such as the IPCC, the Intergovernmental Panel on Climate Change) knows about the causes of climate change. To simplify matters a

[7] Example courtesy of Niko Strobach.

[8] Wayne Riggs diagnoses an asymmetry here: "Theories of epistemic justification can, prima facie at least, emphasize avoiding falsehoods over having truths but not vice versa." (Riggs 2003, p. 247).

[9] Some would like to infer from the observation that (some) theories are underdetermined by the evidence that contextual factors unavoidably intrude into our practice of scientific evaluation. Justin Biddle (2013), for example, argues that in light of transient underdetermination (which, following Kitcher (2001), he defines as the claim that some theories are underdetermined by logic and currently available evidence [cf. Biddle 2013, p. 125)] contextual (i.e. non-epistemic) factors must (and will) be drawn upon in order to decide between competing theories or hypotheses in science.

bit, let the question be whether a particular renowned expert, call her Mira, who is extremely knowledgeable about the topic and has investigated the matter thoroughly, knows that climate change is anthropogenic. We all agree that the consequences of falsely denying knowledge would be socially, economically and morally disastrous. Our distal *practical* (non-epistemic) aim is therefore to avoid these consequences. Our distal *epistemic* aim is to acquire a true belief with respect to the question of whether climate change is anthropogenic. Our proximal, *practico-epistemic* aim is to (not falsely deny knowledge but rather to) certify actionable information—information that one can, and maybe even ought to, act upon. Suppose furthermore that climate change is anthropogenic; suppose also that Mira's evidence is very good though not conclusive—there are still alternative explanations of the data that she cannot completely rule out. Is a knowledge ascription to the effect that Mira knows that climate change is anthropogenic correct in this case? The contextualist would like to answer in the affirmative, given that Mira's epistemic position meets the impure (practico-epistemic) standard that results from weighing epistemic and non-epistemic interests. That is not to deny that whether Mira meets some other, invariant (but also inscrutable) standard, is still an unresolved question; but one the contextualist takes no particular interest in resolving.

Health Take another case. Suppose that you wonder whether to self-attribute knowledge regarding a particular proposition p. Let p be the proposition that a particular substance has no detrimental effect on your health. Again, your distal *epistemic* aim is to acquire a true belief regarding p. Your distal *practical* aim is to lead a long and healthy life and to therefore not consume unhealthy substances. Again, much is a stake; it is important that you get it right. Yet this time the consequences of falsely ascribing knowledge would be disastrous. Your proximal, *practico-epistemic* aim is to (not falsely ascribe knowledge but to still) certify information that you can use as a premise in your practical deliberation. Obviously, you could also refrain from self-ascribing knowledge; you could even deny that you know. The knowledge denial would force you to gather more evidence; it would trigger *epistemic anxiety,* understood as "a desire for increased cognitive activity" (Nagel 2010, p. 414). But then, at one point in the process, you might simply rest content with the evidence gathered so far and think it enough to base your deliberation on. If, at this point, you come to self-ascribe knowledge, the function of the attribution is to certify actionable information and to thereby terminate 'inquiry'. As in many other cases, your epistemic aim of forming a correct doxastic attitude regarding p will have to be balanced against the practical aim of acting in a particular manner: of getting your work done, of publishing a paper, of living a life worth living, etc. Suppose then, that, eventually, after having carefully assessed the evidence, you ascribe to yourself the knowledge that the substance in question has no detrimental effects. Suppose that this is true. Is the ascription correct? Again, the contextualist would like to answer in the affirmative, given that your epistemic position meets the standard set by your practico-epistemic needs. What other, in particular invariant, standard could you be asked to meet in order for the ascription to be correct?

Not surprisingly then, in examining how epistemic and non-epistemic interests and aims interact, a certain function of knowledge ascriptions comes into view: the function of certifying actionable information. That human beings (and organisms

in general) seek information and greatly benefit from tapping others as sources of information is a well-rehearsed point in the literature; the debate owes much to Edward Craig's exposition of the topic (Craig 1990). But they commonly seek information for a reason; and more often than not, this reason seems to be a practical reason. Consequently, it is often *actionable* information that they seek (—an idea that can be found in the writings of David Henderson, Michael Hannon, and others and that I want to elaborate on in what follows).

7 The function(s) of knowledge ascriptions

But why think that certifying actionable information actually *is* the function knowledge ascriptions predominantly serve?[10] Various other candidates for the title "the function of knowledge ascriptions" are currently on offer (cf. Hannon 2015 for an overview); knowledge ascriptions are said to serve the purpose of

(i) flagging good informants (Craig 1987, 1990),
(ii) terminating inquiry or giving discretion over a practice (cf., e.g., Kvanvig 2009b; Kusch 2009; Kelp 2011; Rysiew 2012),
(iii) tracking norms of action or practical deliberation (cf., e.g., Fantl and McGrath 2002; Hawthorne and Stanley 2008),[11]
(iv) keeping epistemic gate (Henderson 2009, 2011),
(v) giving assurance (Lawlor 2013),
(vi) evaluating agents or actions (cf., e.g., McGrath 2015),
(vii) serving as a basis for attributions of blame or moral responsibility (Beebe and Buckwalter 2010; Beebe 2012).

—to name just a few prominent contenders.[12] Are these suggestions in conflict with the idea that knowledge ascriptions predominantly serve the function of certifying actionable information?

(i) Suggesting good informants, it seems to me, is actually not a function of knowledge ascriptions of the form "X knows that p" at all; suggesting good informants is commonly achieved by ascriptions of the form "X knows whether p" which ascribe expertise or epistemic authority regarding a particular subject area. (ii) Secondly, why do we want an inquiry to be stopped? Presumably because we want to

[10] Also, one might wonder whether assertions ("P") not also serve the function of providing actionable information. Yet while they may often do so, knowledge ascriptions ("I know that P") commonly involve a stronger commitment (to the truth of P). They are thus a means not just of providing but of certifying actionable information.

[11] Another much discussed knowledge norm is the knowledge norm of assertion (cf. Williamson 1996). Yet knowledge as a necessary condition for assertion seems too strong, and empirically falsified at that (Kneer 2018).

[12] These considerations are neither meant to contribute to an analysis of the concept of knowledge nor to elucidate the nature of knowledge; the aim is solely to examine the function of *knowledge ascriptions* and their *correctness conditions*.

do something with the results of the inquiry, act upon the insights thus garnered. If that is so, what better way is there to legitimately stop an inquiry than by providing actionable information, information that can henceforth be used as premise in practical (and also further theoretical) deliberation; that can be acted upon? Similarly, why do we give another person discretion over a practice (—such as mushroom hunting; Kusch 2009, p. 84)? Presumably because we believe that the other person is more 'knowledgeable' in relevant respects and can give us the information we need in order to successfully engage in that practice ("She knows that this mushroom is edible"). (iii) Certain 'norms' of action will thereby be tracked, as certain actions will become permissible, other things being equal. (iv) In certifying actionable information, knowledge ascriptions also serve the aim of keeping epistemic gate (Henderson 2011). (v) Moreover, giving assurance serves the purpose of providing the addresses with exclusionary reasons (Lawlor 2013, p. 17); reasons, i.e., that theoretical as well as practical deliberation can be based on. —What about the last two suggestions: (vi) agent evaluation and ascribing knowledge in order to justify attributions of blame (vii)? The latter might be taken to be a special case of the former. In the former, "we criticize and defend actions by attributing knowledge" (McGrath 2015, p. 138), as in the following example:

> "The tub has a drip. Your spouse is applying great force to the cold-water handle. You finally say, 'You know that's not going to work! Let's call the plumber'" (McGrath 2015, p. 138).

In agent evaluation, too, we are certifying information as actionable: as something that the subject ought to act or have acted upon. What your spouse is actually saying is: "You know that's not going to work—so act accordingly!".

So, unless one thinks that a knowledge ascription fulfills no function at all, certifying actionable information is the most promising candidate for the function it fulfills. But even if we agree that certifying actionable information is an important function of knowledge ascriptions, further questions still await an answer. One question pending is what exactly 'actionable' is supposed to mean. Different readings need to be distinguished. The function of a knowledge ascription may be either

(i) to certify or provide information that *ought* to be acted upon; or.
(ii) to certify or provide information upon which one is *permitted* to base one's practical deliberation.

The latter formulation, again, permits of two different readings;

(a) as providing *exclusionary reasons* (e.g., Lawlor 2013, p. 17); reasons that one can act upon without further consideration, or.
(b) as providing a sort of *prima facie reason* to act, i.e., as providing information that is fit to serve as a premise in practical deliberation but still has to be balanced against other reasons or information.

These considerations bear on the question of how intimate a connection there is between knowledge ascriptions and action. Knowledge has, occasionally, been said to be necessary for action; it has also been said to be sufficient for action. Here are two versions of these claims:

> "NEC: if it is appropriate to rely on p in practical reasoning then you know that p.
> SUFF: if you know that p then it is appropriate to rely on p in practical reasoning" (Brown 2008, p. 168).

Various arguments against either NEC or SUFF have been put forward. I won't attempt to defend NEC as I take it to be (and many others have argued that it is) too strong a claim anyway (cf. Baumann 2012 for an interesting alternative). But it might seem as if I were committed to SUFF. Given that what is at issue here are knowledge ascriptions, and in light of the distinctions just drawn, let us reformulate SUFF accordingly:

> $SUFF_{KA}$: If by means of a knowledge ascription information is certified as actionable for an agent A, then A either
>
> (i) ought to act upon the information thus certified (strong reading); or.
> (ii) is permitted to base his or her practical deliberation on the information thus certified (weak reading), by being given either.
>
> (a) exclusionary reasons (weak $reading_{ER}$), or
> (b) prima facie reasons (weak $reading_{PFR}$).

Note that this holds good only if the ascription (qua speech act) was felicitous; and only if the function of the ascription indeed was to certify actionable information. As I am not claiming that this is the *only* function of knowledge ascriptions, let alone always fulfilled, SUFF needs to be qualified accordingly. Also, I favor the weak $reading_{PFR}$ (reading (ii)b). I am therefore committed only to the following: If a knowledge ascription is felicitous and fulfills the function of certifying actionable information (for A), then A has thereby been given a prima facie reason to act upon the information thus certified, i.e. A has been provided with information that can serve as a premise in practical deliberation but still has to be balanced against other reasons or information. A has not been given a license to act or an all-things-considered reason to act. In particular, A might still have other (e.g., moral) reasons not to act.

8 Objections and replies

Objection 1 Might not someone pursue purely epistemic aims after all? And wouldn't a respective knowledge ascription thereby invoke a 'pure' epistemic standard?[13]

Reply Someone might indeed pursue purely epistemic aims. In certain cases, which are, presumably, rather recherché and purely theoretical, we might try to ignore practical concerns altogether, as Descartes requested of the readers of his *Meditationes*. This will land us with a purely epistemic standard. However, one ought to distinguish different cases, depending on which epistemic aim is to be pursued.

a. One might try to pursue both aims and maximize true beliefs while at the same time minimizing false beliefs. As discussed before, a creature with no practical interests might do so without compromise. Being neither pressed for time nor afflicted by any other practical concerns, it might go about acquiring one error-proof belief after another. Acting creatures like us, on the other hand, being motivated and constrained by practical concerns and factors, have to trade in quality for quantity or vice versa.
b. One might instead pursue the sole aim of minimizing false beliefs, and reserve the term 'knowledge' for those beliefs which are error-proof. That will result in a skeptical standard.
c. Or one might pursue the sole aim of maximizing true beliefs. What epistemic standard would that result in? Interestingly, the term 'epistemic' is used to either denote something like truth-conducive factors or something that turns true belief into knowledge (cf. Stanley 2005, p. 2). Epistemic standards, recall, are those standards someone has to live up to in order to count as a knower (in a particular context of ascription). Consequently, the 'epistemic' in 'epistemic standard' is to be read in the second, knowledge-related way. Now, if one were to maximize true beliefs by simply forming as many beliefs as possible, assuming that the more beliefs one had, the more true beliefs would be among them (and not caring whether one thereby also acquired many false beliefs), that would not result in an epistemic standard at all. Lucky guessing is not the same as knowing. Some attempt at avoiding error has to be made, is seems. Thus, any non-skeptical epistemic standard worth its name will result from balancing the aim of maximizing true beliefs against the aim of minimizing false beliefs. Still, the standard operative in a particular context is not set in an arbitrary manner, as it is the standard that squares the agents' epistemic (e.g., the amount of certainty or robustness they require) and non-epistemic needs—or so I suggest.

[13] I am very grateful to an anonymous referree for asking me to clarify this point.

Objection 2 This raises a related objection. For couldn't differences in epistemic standards also result from weighing different epistemic aims differently? More specifically, couldn't they be weighed differently on purely epistemic grounds?

Reply The aim of maximizing true beliefs could be balanced against the aim of minimizing false beliefs in different ways. Yet unless there is a further, overarching epistemic aim (something I don't want to rule out, but that I am not prepared to argue for here either) that could be decisive, the reasons for balancing them in one way rather than another have to be non-epistemic reasons, it seems. But then, other cases are also possible (as has been pointed out to me by an anonymous referee). For example, the aim of maximizing true beliefs across domains could be balanced against the aim of maximizing true beliefs in a particular domain, leading to different epistemic standards. Yet again, at least some of the reasons for choosing between the two different epistemic aims will stem from practical interests and concerns; for, which purely epistemic (i.e. either truth-related or knowledge-related) reasons could decide between the two cases? I can think of none (lack of imagination is, of course, not an argument, but it is the best I can do here).

Objection 3 Someone might object that there is an invariant yet communally determined standard after all, the standard that is applied in science. For isn't the scientific standard, arguably, thought to meet the demands of all comers and to certify information that can be drawn on for all kinds of purposes? Even a cursory glance at scientific papers (by neuropsychologists or biologists, for example) lends support to this conjecture, as they typically make little to no knowledge ascriptions. Authors of such papers say things like: "these findings suggest…", "these findings indicate…", "we propose…", "our work supports…" or, at most, "our results show…".[14] This can be nicely explained (admittedly, a nice explanation need not be a correct explanation) on the assumption that they take a very high standard to be operative in their respective contexts; and they try (and ought) to take the interests of all potential addressees into account. Yet the less clear it is, who is going to draw on the information provided, the more potential interests need to be accommodated. Standards will be raised accordingly. I therefore venture the guess, and put forth as an empirical hypothesis (that a corpus-linguistics study, for example, might investigate), that the less clear it is who is going to draw on a piece of information to be conveyed and for what purpose this is done, the less inclined (epistemically vigilant) speakers are to make knowledge claims. In broader-purpose or general-purpose source contexts (— as opposed to "applied practical communities", Henderson and Horgan 2015, p. 99), people tend to make fewer knowledge ascriptions, or so I conjecture. The scientific standard, therefore, may be the most plausible candidate for a communally determined standard (cf. also Henderson and Horgan 2015, p. 95).

Reply Yet one might doubt that there *is* a single scientific standard that remains *invariant* across time and scientific communities. Also, if there were, it would be

[14] This may hold good for claims made about 'cutting edge' research, but less so for claims made about established scientific results (I am grateful to Charles Lowe for alerting me to the need to draw this distinction). Yet this could be explained on the assumption that established results have stood the test of time and so have, presumably, been already (successfully) drawn on by quite divers comers.

a very demanding standard and warrant only very few knowledge ascriptions. The more (and diverse) comers it would have to satisfy, the more demanding it should become. It would not be the standard for everyday knowledge. I thus take the idea of a single communal, invariant yet *moderate* standard to be as attractive as it is problematic. And I surmise that ascribing knowledge is commonly a more 'local', context- or task-sensitive endeavor.

Objection 4a Jessica Brown objects to SUFF. Since I am committing myself to a version of SUFF (if only a weak one), I should comment on her argument. According to Brown (2008), cases of bets with very odd stakes provide counterexamples to SUFF, as when Liz, who is pretty sure that she was born in England, is asked to accept the following bet: she gains £1 if she was born in England but loses her home if she was not born in England. Brown describes the case as follows: Liz knows that she was born in England; yet she should not accept the bet (ibid., p. 176).

Reply a But then, might not one describe the case somewhat differently and claim that the fact that she should not accept the bet is a reason (due to high stakes) not to ascribe her knowledge about her place of birth in the first place, exactly because the ascription would not provide her with actionable information?

Objection 4b Another counterexample is the case of the surgeon, who before operating on a diseased kidney, double-checks the patient's records to make sure that he is operating on the kidney that is in fact diseased, and not the healthy one (Brown 2008, p. 176; cf. Hannon 2015, p. 860). While he knows which kidney to operate on, or so Brown has it, he should not act on that knowledge but double-check first.

Reply b This raises what McGrath dubs 'the epistemic objection', for doesn't it seem odd to say "The surgeon knows which kidney to operate on but he ought not to/cannot act on that knowledge" (McGrath 2015)? Moreover, suppose he has double-checked. What epistemic status has he thereby earned? Shall we ascribe him super-knowledge? But might not someone object that although he super-knows which kidney to operate on he should nonetheless triple-check, given what is at stake. What if he has triple-checked, and we ascribe him super-duper-knowledge; mightn't the same objection be raised again and again, thereby exploiting the infallibilist intuition that one can never be sure unless one can rule out all error possibilities (McGrath 2015)? If knowledge doesn't license action, then neither does super-knowledge, nor any other epistemic state short of infallibilist knowledge. But what function does the concept of knowledge then serve? Ascribing knowledge may be a way of blocking the infallibilist objection; it may (often) be a way of suggesting that "enough is enough", as Austin puts it (1946, p. 84).

9 Whose stakes?

There is also another question that needs to be addressed, namely the question of *whom* the information is supposed to be actionable for: the speaker (SP), who makes the ascription; the addressee (AD), whom the speaker addresses in making the ascription; or the putative knower, who is the subject of the ascription (SU)?

(1) Suppose *A* says to *B*, who wonders when the opera begins:

"I know that the Opera begins at 8 o'clock".

A is giving assurance and, arguably, provides *B*, the *addressee* (AD), with actionable information as *A* is implying that *B* can act upon the information thus provided. Stakes come into play, as the more is at stake for *B*, the better *A's* reasons or evidence have to be.

(2) Suppose *A* has second thoughts herself and wonders whether the opera begins at eight. After having checked her sources, she concludes:

"I know that the Opera begins at 8 o'clock".

In this case, *A* is providing *herself*, who happens to also be the *addressee* and the *subject* of the ascription (AD = SP = SU), with actionable information; her evidence is good enough given her practical interests.

(3) Suppose *A* says to her partner, *B*, who is not yet dressed:

"You know that the Opera begins at 8 o'clock".

This is a case of agent evaluation, and *A* seems to be suggesting that *B*, the *addressee,* who is also the *subject* of the ascription (AD = SU), ought to act upon the information thus certified.

(4) Suppose *A* says to *B*, who wonders whether *C* will be on time:

"*C* knows that the opera begins at 8 o'clock".

This is a case of action explanation or prediction; *A* is implying that *C* has the relevant piece of information and will act accordingly. In this case, the information is supposed to be actionable for a *third party*, *C*, who is the *subject* of the ascription (SU = TP).

(5) Finally, suppose *B* is asking when the opera begins and *A* answers as before:

"*C* knows that the opera begins at 8 o'clock".

C might have told *A* so or might have provided her with evidence; or *A* might simply trust *C* to be well informed about the matter at hand. *A* is again providing *B*, the *addressee* (AD), with actionable information.

In some cases, therefore, the ascriber is certifying information as being 'actionable' for the addressee, who might be identical to the speaker or to the subject of the ascription. In other cases, she is certifying information as being actionable for the subject of the ascription (the putative knower), who might be identical to the speaker, the addressee or a third party. Not surprisingly then, and as Matthew

McGrath notes, Subject Sensitive Invariantism (SSI) is well suited to accommodate those cases in which information is certified as actionable for the subject, as what determines the epistemic standard according to SSI are the practical interests and stakes of the subject (McGrath 2015). Contextualism, on the other hand, is better suited to accommodate cases in which information is certified as actionable for the addressee, as in these cases their interests ought to determine the operative standard. One might speculate that defenders of SSI mostly had cases of agent evaluation in mind when developing their theory, contextualists mostly cases of certifying information for an addressee; both can be faulted for a biased diet of samples, it seems.

The lesson to be learned, then, is that the different functions knowledge ascriptions may fulfill require that different sets of interests and needs be taken into account. They in turn set the standard that is operative in the context of ascription. If the function of a particular knowledge ascription is to provide information that is actionable for the addressee, her interests and practico-epistemic needs (or stakes) count and the standard operative in the context of ascription is set accordingly. If the function is to certify information as actionable for the subject, her interests (or stakes) count. If the ascriber is mistaken about the addressee's or the subject's interests, then the knowledge ascription 'misfires', to use Austin's phrase. But then, although ascribing knowledge may, often, be a local (dyadic or triadic) undertaking, requiring that the ascriber be sensitive to someone else's needs, there may also be community norms governing knowledge ascription, especially if it is to be expected that the information (thus provided) will travel around.[15] Also, as discussed before, people occasionally ascribe knowledge with no *particular* interest in mind. They broadcast information throughout the community, information that is supposed to be good for all (or different) kinds of purposes and to meet the demands of all (or different types of) comers.

As Keith DeRose has pointed out, it is compatible with (ascriber) contextualism that the ascriber selects a standard as she sees fit (DeRose 2005, p. 189), thereby going flexible, as Crispin Wright aptly puts it (Wright 2017).[16] She may select standards that are "appropriate to the subject's practical situation" or standards "different from those that the subject's own practical situation calls for" (DeRose 2005, p. 189). Elaborating on this idea, the account developed in this paper suggests that the standard is selected in light of the function the ascription is supposed to fulfill. The function will reflect the epistemic needs of the addressee or the subject of the ascription. The correctness-condition for knowledge ascriptions needs to be amended accordingly:

> NIC*: A knowledge ascription of the form "A knows that P" in context c is correct only if the proposition thereby expressed is true relative to circumstances (w_c, e_c), where w_c is the world of c and e_c is the epistemic standard operative in c. Which standard is operative in a given context is determined

[15] I am grateful to an anonymous referee for helpful comments on this point.

[16] As mentioned before, this is similar to Krista Lawlor's point that the reasonable person standard is set by what "a reasonable person with the speaker's interest in the question of p's truth and reasonable expectations about the hearer's interests in this question" (Lawlor 2013, p. 160) would take into account.

by the function the ascription is supposed to fulfill. If it is supposed to certify information as actionable for the subject, the subject's practical interests count (i.e. set the standard); if it is supposed to certify information as actionable for the addressee, his or her practical interests count. If it is supposed to provide information for all kinds of agents and purposes, all their interests count, resulting in a potentially extremely high standard (operative in the context of ascription) that respects the needs of all those who could possibly draw on the information thus provided.

10 Summing up

I have been arguing that practical interests come into play when agents decide which practico-epistemic goal to pursue in a particular situation. Commonly, agents seek actionable information. One main purpose of knowledge ascriptions is, therefore, to certify information as actionable. Consequently, the practical (and epistemic) interests of those for whom the information is supposed to be actionable help setting the standard operative in the context of ascription. The invariantist, on the other hand, owes us an explanation of what function or purpose knowledge ascriptions 'evolved' to serve. As Henderson puts it:

> "Plausibly, evaluative concepts arose with an eye to, or sensitivity to, what is needed, or conduces to, successful pursuit of some project" (Henderson 2011, p. 86).

The invariantist needs to come up with an alternative function of knowledge ascriptions. He might claim that they are to pick out particularly good epistemic positions. But the interesting question is: why do we care about people in good epistemic positions? Moreover, how does the function fit in with the semantics? Words acquire meaning by being used in particular ways; if usage changes, so does meaning. It is incumbent upon the invariantist to explain how speakers' linguistic behavior manages to set an invariant threshold for knowledge (or how else the word 'know' acquires its meaning), and why people should care about it.

Acknowledgements The paper grew out of a joint research project with Igal Kvart (Hebrew University). I gratefully acknowledge financial support from the State of Lower Saxony, Hannover, Germany and the Volkswagen Foundation (VWZN3163). I am indebted to Igal Kvart, who invited me to Jerusalem to present an earlier version of the manuscript, for fruitful discussions. I would also like to thank the participants of the conference "The epistemic significance of non-epistemic factors" (funded by the VW Foundation and the DFG, KO 1971), which Igal Kvart, Andrea Robitzsch and I jointly organized in July 2018, especially Krista Lawlor, Sandy Goldberg, and David Henderson. Furthermore, I am very grateful to other audiences at the universities of Münster and Osnabrück, in particular to Susanne Boshammer, Charles Lowe, Christian Nimtz, Sebastian Schmoranzer, and Niko Strobach, for valuable comments and very helpful feedback. My greatest debt goes to Andrea Robitzsch for reading and commenting on earlier drafts of the manuscript and for many an inspiring discussion. Also, I would like to thank Charles Lowe again for proof-reading the whole manuscript. Finally, I would like to thank two anonymous referees for this journal for their comments and helpful suggestions.

Funding Open Access funding enabled and organized by Projekt DEAL.

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

Annis, D. (1978). A contextualist theory of epistemic justification. *American Philosophical Quarterly, 15,* 213–219.

Austin, J. L. (1946). Other minds. *Proceedings of the Aristotelian Society, 148,* 148–187. (Reprinted in J. L. Austin ^{3}1979. *Philosophical Papers* (pp. 76–116). Oxford: Oxford University Press).

Bach, K. (2006). The excluded middle: Semantic minimalism without minimal propositions. *Philosophy and Phenomenological Research, 73*(2), 435–442.

Baumann, P. (2012). Knowledge, practical reasoning and action. *Logos and Episteme,* III(I), 7–26.

Baumann, P. (2016). *Epistemic contextualism—A defense*. Oxford: Oxford University Press.

Beebe, J. (2012). Social functions of knowledge attributions. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions* (pp. 220–242). Oxford: Oxford University Press.

Beebe, J., & Buckwalter, W. (2010). The epistemic side-effect effect. *Mind and Language, 25*(4), 474–498.

Biddle, J. (2013). State of the field: Transient underdetermination and values in science. *Studies in History and Philosophy of Science, 44,* 124–133.

Blome-Tillmann, M. (2014). *Knowledge and presuppositions*. Oxford: Oxford University Press.

Brown, J. (2006). Contextualism and warranted assertability manoeuvres. *Philosophical Studies, 130*(3), 407–435.

Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Nous, 42*(2), 167–189.

Carston, R. (2002). *Thoughts and utterances: The pragmatics of explicit communication*. Oxford: Blackwell Publishing.

Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. In J. E. Tomberlin (Ed.), *Philosophical perspectives* (Vols. 13, pp. 57–89). Oxford: Blackwell.

Craig, E. (1987). The practical explication of knowledge. *Proceedings of the Aristotelian Society, 87,* 211–226.

Craig, E. (1990). *Knowledge and the state of nature: An essay in conceptual synthesis*. Oxford: Oxford University Press.

Davis, W. A. (2004). Are knowledge claims indexical? *Erkenntnis, 61,* 257–281.

Davis, W. A. (2005). Contextualist theories of knowledge. *Acta Analytica, 20,* 29–42.

Davis, W. A. (2007). Knowledge claims and context: Loose use. *Philosophical Studies, 132,* 395–438.

Davis, W. A. (2017). Loose use and belief variation. In J. J. Ichikawa (Ed.), *The Routledge handbook of epistemic contextualism* (pp. 218–229). London, New York: Routledge.

DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research, 52*(4), 913–929.

DeRose, K. (1999). Contextualism: An explanation and defense. In J. Greco & E. Sosa (Eds.), *The Blackwell guide to epistemology* (pp. 187–205). Oxford: Blackwell Publishers.

DeRose, K. (2000). Now you know it, now you don't. *Proceedings of the Twentieth World Congress of Philosophy, 5,* 91–106.

DeRose, K. (2005). The ordinary language basis for contextualism and the new invariantism. *The Philosophical Quarterly, 55,* 172–198.

Douglas, H. (2000). Inductive risk and values in science. *Philosophy of Science, 67*(4), 559–579.

Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *The Philosophical Review, 111,* 67–94.
Gerken, M. (2012). On the cognitive basis of knowledge ascriptions. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions* (pp. 140–170). Oxford: Oxford University Press.
Gerken, M. (2013). Epistemic focal bias. *Australasian Journal of Philosophy, 91*(1), 41–61.
Grimm, S. R. (2015). Knowledge, practical interests, and rising tides. In D. Henderson & J. Greco (Eds.), *Epistemic Evaluation* (pp. 117–137). Oxford: Oxford University Press.
Hannon, M. (2015). The importance of knowledge ascriptions. *Philosophy Compass, 10*(12), 856–866.
Hannon, M. (2017). A solution to knowledge's threshold problem. *Philosophical Studies, 174*(3), 607–629.
Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy, 105*(10), 571–590. [52(4), 919].
Henderson, D. (2009). Motivated contextualism. *Philosophical Studies, 142,* 119–131.
Henderson, D. (2011). Gate-keeping contextualism. *Episteme, 8*(1), 83–98.
Henderson, D. (2020). The place of non-epistemic matters in epistemology: Norms and regulation in various communities. *Synthese* (online first). https://doi.org/10.1007/s11229-020-02935-2.
Henderson, D., & Horgan, T. (2015). What's the point? In D. Henderson & J. Greco (Eds.), *Epistemic evaluation* (pp. 87–114). Oxford: Oxford University Press.
Kaplan, D. (1989). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford: Oxford University Press.
Kelp, C. (2011). What's the point of "knowledge" anyway? *Episteme, 8*(1), 53–66.
Kitcher, P. (2001). *Science, truth, and democracy*. Oxford: Oxford University Press.
Kneer, M. (2018). The norm of assertion: Empirical data. *Cognition, 177,* 165–171.
Kölbel, M. (2008). Introduction: Motivations for relativism. In M. García-Carpintero & M. Kölbel (Eds.), *Relative truth* (pp. 1–38). Oxford: Oxford University Press.
Kompa, N. (2002). The context sensitivity of knowledge ascriptions. *Grazer Philosophische Studien, 64,* 1–18.
Kompa, N. (2005). The semantics of knowledge attributions. *Acta Analytica, 20*(1), 16–28.
Kompa, N. (2012). Nonindexical contextualism—An explication and defence. In S. Tolksdorf (Ed.), *Conceptions of Knowledge* (pp. 329–349). Berlin: DeGruyter.
Kompa, N. (2014). Knowledge in context. *Rivista Internazionale di Filosofia e Psicologia, 5,* 58–71.
Kompa, N. (2015). Contextualism and disagreement. *Erkenntnis, 80*(1), 137–152.
Kompa, N. (2016). Faultless disagreement, context sensitivity, and the semantics of evaluation. *Grazer Philosophische Studien, 93,* 396–418.
Kompa, N. (2017). Meaning and interpretation. In S.-J. Conrad & K. Petrus (Eds.), *Meaning, context, and methodology* (pp. 75–90). Berlin/New York: De Gruyter.
Kusch, M. (2009). Testimony and the value of knowledge. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic value* (pp. 60–94). Oxford: Oxford University Press.
Kvanvig, J. (2009a). The value of understanding. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic value* (pp. 95–111). Oxford: Oxford University Press.
Kvanvig, J. (2009b). Responses to critics. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic value* (pp. 339–351). Oxford: Oxford University Press.
Kvart, I. (2018). Pragmatic structures for action-directed pragmatics. *Philosophical Perspectives, 32*(1), 219–253.
Lawlor, K. (2013). *Assurance. An Austinian view of knowledge and knowledge claims*. Oxford: Oxford University Press.
Lewis, D. (1980). Index, context, and content. In S. Kangar & S. Öhman (Eds.), *Philosophy and grammar*. Dordrecht: Reidel. (Reprinted in D. Lewis (1998) *Papers in Philosophical Logic* (pp. 21–44). Cambridge: Cambridge University Press).
MacFarlane, J. (2005). The assessment sensitivity of knowledge attributions. In T. Szabo Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (pp. 197–233). Oxford: Oxford University Press.
MacFarlane, J. (2007). Semantic minimalism and nonindexical contextualism. In G. Preyer & G. Peter (Eds.), *Context-sensitivity and semantic minimalism: New essays on semantics and pragmatics* (pp. 240–250). Oxford: Oxford University Press.
MacFarlane, J. (2009). Nonindexical contextualism. *Synthese, 166,* 231–250.
MacFarlane, J. (2014). *Assessment sensitivity: Relative truth and its applications*. Oxford: Oxford University Press.

McGrath, M. (2015). Two purposes of knowledge attribution and the contextualism debate. In D. Henderson & J. Greco (Eds.), *Epistemic evaluation* (pp. 138–161). Oxford: Oxford University Press.
McKenna, R. (2013). 'Knowledge' ascriptions, social roles and semantics. *Episteme, 10*(4), 335–350.
Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy, 86*(2), 279–294.
Nagel, J. (2010). Epistemic anxiety and adaptive invariantism. *Philosophical Perspectives, 24,* 407–435.
Nagel, J., & Smith, J. J. (2017). The psychological context of contextualism. In J. J. Ichikawa (Ed.), *The Routledge handbook of epistemic contextualism* (pp. 94–104). London and New York: Routledge.
Récanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.
Récanati, F. (2008). Moderate relativism. In M. García-Carpintero & M. Kölbel (Eds.), *Relative truth* (pp. 41–62). Oxford: Oxford University Press.
Riggs, W. (2003). Balancing our epistemic goals. *Noûs, 37*(2), 342–352.
Robitzsch, A. (2019). *An externalist approach to epistemic responsibility: Intellectual norms and their application to epistemic peer disagreement (Synthese Library, 411) (Synthese Library, 411)*. New York: Springer.
Rysiew, P. (2007). Speaking of knowing. *Nous, 41*(4), 627–662.
Rysiew, P. (2012). Epistemic scorekeeping. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions* (pp. 270–293). Oxford: Oxford University Press.
Rysiew, P. (2017). 'Knowledge' and pragmatics. In J. J. Ichikawa (Ed.), *The Routledge handbook of epistemic contextualism* (pp. 205–2017). London and New York: Routledge.
Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
Stei, E. (2014). *Die Bedeutung von "Wissen". Eine Untersuchung zur Kontextabhängigkeit von Wissensaussagen.* Münster: mentis.
Williamson, T. (1996). Knowing and asserting. *Philosophical Review, 105*(4), 489–523.
Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *The Philosophical Quarterly, 55*(219), 213–235.
Crispin, W. (2017). The variability of 'knows': An opinionated overview. In J. J. Ichikawa (Ed.), *The Routledge Handbook of Epistemic Contextualism* (pp. 13–31). London, New York: Routledge.

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