THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

From Epistemic expressivism to epistemic inferentialism

Citation for published version:

Chrisman, M 2010, From Epistemic expressivism to epistemic inferentialism. in A Haddock, A Millar & D Pritchard (eds), *Social Epistemology*. Oxford University Press.
<https://doi.org/10.1093/acprof:oso/9780199577477.003.0006>

Digital Object Identifier (DOI):

[10.1093/acprof:oso/9780199577477.003.0006](https://doi.org/10.1093/acprof:oso/9780199577477.003.0006)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Published In:

Social Epistemology

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

From Epistemic Expressivism to Epistemic Inferentialism

Matthew Chrisman

forthcoming in *Social Epistemology*, Haddock, Pritchard, Millar (eds.), Oxford University Press

1. INTRODUCTION

Ordinary people often make epistemic claims. For example, they claim to know the whereabouts of a misplaced book, or that someone else knows whether Jones will be at the party, and they might even do something reconstructable as claiming that Suzy's belief is justified even though false or claiming that Henry's belief is justified and true though not a case of knowledge. Insofar as they make such claims, presumably they are at least somewhat concerned to do so rightly. Epistemologists make such epistemic claims too and are, of course, also concerned to do so rightly. However, they often give priority over these epistemological issues to what I think can only be viewed as meta-epistemological issues. These are issues not about whether this or that epistemic claim is correct but about things like the standard for knowledge, the nature of epistemic justification, the possibility of rational response to skeptical arguments, and the meaning of epistemic claims.

Concerning this final issue, recent philosophical debate has largely centered on the question of whether epistemic claims are plausibly thought to be context sensitive. The default assumption has been that sentences that attribute knowledge or justification (or whatever else is epistemic) have stable truth-conditions across different contexts of utterance, once any non-epistemic context sensitivity has been fixed. The idea, for instance, is that sentences of the form 'S knows that p' have stable truth-conditions once the values of S and p (plus any implicit temporal elements) are fixed. The contrary view is the contextualist thesis that such sentences do not have stable truth-conditions but can vary depending on the context of utterance. So, for instance, even once we've fixed the values of S and p the sentence 'S knows that p' could still be true when tokened in one context, while false when tokened in another context, since the truth-conditions of this sentence depend in part on some contextually varying element such as the strength of evidence required for knowledge.¹

This debate manifestly presupposes that the meta-epistemological issue of accounting for the meaning of epistemic claims is to be settled by determining the truth-conditions of these claims. I think this presupposition is probably false and, in any case, way too quick. This is for two interlocking reasons. First, many epistemologists see epistemic claims as evaluative or normative, in some sense.² However, in the meta-*ethical*

¹ See for instance Unger (1984), Cohen (1988), DeRose (1992, 1995, 2009), Lewis (1996), Rysiew (2001), Neta (2003), Hawthorne (2004), Stanley (2005).

² See for instance the virtue epistemology of Sosa (1980; 2007) and Zagzebski (1996), but also those who think knowledge is a norm of belief, action, and/or assertion, e.g. Pollock and Cruz (1999: ch. 5),

debate – which has really become a meta-*normative* debate – most participants take alternatives to truth-conditional semantics, such as expressivism, as live options. Second, from a pragmatist view on meaning (that I think makes contact with the remit of ‘social epistemology’) a natural way to approach the meaning of epistemic claims would be to ask not about their truth-conditions but about what practical function they serve in our commerce with one another and other aspects of the world. It could be that, even if we agree about the truth-conditions of knowledge attributions, there’s still an important meta-epistemological question to be settled about their linguistic and conceptual role.

These two reasons correlate with two moments in a line of thought whose name I have taken for the title of this paper. My purpose in it is to flesh out this line of thought and then indicate how I think it could help to enrich meta-epistemological theorizing in a way that appreciates the specially social character of epistemic claim-making.

2. ETHICAL EXPRESSIVISM AND EPISTEMIC EXPRESSIVISM

In the meta-ethical debate, ethical expressivists encourage us to ask not about the nature of ethical facts or values but instead about the nature of ethical evaluations. They do so because they think we can satisfactorily answer the latter question without ever countenancing some *sui generis* kind of thing – ethical facts or values – to answer the first question. Their idea is that ethical evaluations are not representations of how the world is; rather, they are expressions of our attitudes towards the way we represent the world to be. Specifically, ethical claims are seen not as descriptions of the world but as expressions of our ethical attitudes towards things like actions, agents’ characters, systems of rules, social structures, etc.

From the perspective of a truth-conditional semantics, which takes notions like truth and reference as its master concepts, the semantics for ethical claims that results from the standard sort of expressivism will appear quite radical.³ In effect, it is the view that ethical sentences mean what they do in virtue of the state of mind that they express, rather than in virtue of the facts they purport to represent.

However, this semantical idea is not without at least some historical inspiration and contemporary allies in the philosophy of language. A rich alternative semantic program, which is at home with this idea across the board, traces back through Grice at least to Locke. This is the ‘ideationalist’ or ‘psychologistic’ approach to meaning, which is inspired by Locke’s view that ‘Words in their primary or immediate Signification, stand for nothing, but the Ideas in the Mind of him that uses them’(1689: 405). In its contemporary manifestation, the project is to understand semantic content in terms of conventional regularities linking overt use of a sentence with a particular idea or state of mind.⁴ Accordingly, all sentences are thought to mean what they do in virtue of the state of mind that they express. And, if we embrace this idea in our philosophy of language, it then opens up room for the expressivist to argue that the conventional regularities of our language are such that descriptive sentences mean what they do in virtue of expressing specific beliefs while ethical sentences mean what they do in virtue of expressing specific

Williamson (2001), Hawthorne and Stanley (2008), and those who follow Sellars (1956) in thinking that knowledge attributions ‘place one in the space of reasons’, e.g. Williams (1992; 2001), Brandom (2000: ch. 2) and Rosenberg (2002).

³ In Bar-On and Chrisman (2009) and Chrisman (2009: §5), nonstandard forms of expressivism are explored, which attempt to be semantically neutral. However, in this paper about the meaning of normative claims, such as epistemic claims, I set those views aside as not addressed to the present issue.

⁴ See Grice (1957), Schiffer (1972), Davis (2003).

desire-like moral attitudes. Because of this, I think expressivism is naturally married to an ideationalist alternative to truth-conditional semantics.⁵

But what exactly does the expressivist say is the desire-like moral attitude expressed by ethical sentences? The details won't matter much here, but to have a toy expressivist theory on the table, we can use a simplification of Gibbard's (1990) leading idea. That is, we can think of the sentence 'X is good' as meaning what it does in virtue of the fact that it express one's acceptance of a system of ethical norms which one believes to encourage the promotion of x.⁶ Expressivist views like this one purport to carry at least two theoretical advantages, which I'll now try to characterize in a rough and ready way.

First, expressivist views claim to have a very simple and intuitive explanation of the apparent practicality of ethical discourse. The rough idea is that the making of ethical claims seems to be intimately and causally linked with motivation to action; if, for example, one judges that giving to charity is good but displays no motivation to give to charity, then something seems to be amiss. And on a popular Humean view⁷ motivation to action always requires the agent to have some desire-like attitude. Expressivism gives us just such an attitude in its account of what is expressed by ethical claims, which generates its explanation of the intuitive practicality of ethical claim-making.

Second, expressivist views claim to avoid queerness/placement problems for ethical facts and values. This is initially an ontological issue, but it is closely related to a conceptual issue having to do with a different sense in which ethical discourse is intuitively practical. The initial idea is that ethical facts, conceived as judgment-independent facts in the world, seem to be neither ontologically reducible to other facts whose existence is not in question nor locatable in the spatio-temporal world governed by empirically discoverable physical laws. Why? Because ethical judgments seem to have a special connection to motivation to action. Roughly, they seem to commit their author to a reason to act in a particular way, irrespectively of any desire to do so. For example, if one judges that giving to charity is good, then one seems to be committed to thinking that there is a reason to give to charity. If we view such judgments as representing the world, then they purport to describe facts which, if they existed would seem to have the power to attract our wills irrespectively of our desires. This can make ethical facts appear queer in the sense that they purport to have a reason-giving power not had by any other kind of fact.⁸ Now, there is plenty of room for debate about whether this is the only plausible conception of ethical concepts and facts or whether the reduction of ethical facts to less queer facts is after all possible. However, because of the very controversies involved in this debate, any meta-ethical view ontologically committed to ethical facts or values bears a theoretical cost not borne by views that do not. Expressivism claims to

⁵ Compare Schroeder (2008) and Chrisman (ms) for more discussion of the semantic commitments of expressivism. Early expressivists like Ayer and Carnap can be plausibly read as denying that ethical terms have semantic value at all. However, this has come to seem increasingly implausible, which is why most recent expressivists should be read, I believe, as relying on an ideationalism as their general semantic approach.

⁶ His is, at its core, a theory of judgments about *rationality* rather than a theory of judgments about something being *good*. I'm simplifying for present purposes. What he actually writes is this: 'An observer, we may assume, accepts at least an incomplete system of general norms; call that system *N*. He then thinks an alternative rational if and only if he thinks it *N-permitted*—that is to say, rational according to *N*...whether an alternative is *N-permitted* is a matter of fact, and so a belief that the alternative is *N-permitted* is a factual belief. Thinking X rational, then, is a combination of a normative state and a state of factual belief'(1990: 91)

⁷ See especially Davidson (1980: 5), Smith (1987; 1990: ch. 4).

⁸ See Mackie (1977: ch. 1).

avoid this cost in its contention that ethical concepts (and perhaps normative concepts more broadly) have their distinctive function in making claims that express desire-like attitudes rather than describe or represent the world.⁹

That's a brief gesture at what ethical expressivism is and two reasons for endorsing it. It's a curious fact that very few philosophers have sought to defend or even recognized the theoretical possibility of a meta-epistemological analogue to ethical expressivism.¹⁰ This is in spite of several obvious parallels: First, it's now commonplace (even if still controversial) for philosophers to insist that epistemic notions like 'knowledge' and 'justification' are normative, and many epistemologists are interested in norms of belief.¹¹ Second, some epistemologists have become quite puzzled about the nature and location of epistemic values, thinking now that if the value of knowledge were merely instrumental value for the goal of truth that it would be 'swamped' once one has the truth by whatever means.¹² Third, several pragmatist-inspired philosophers have asked about the social role of epistemic concepts such as 'knowledge' and come up with an account according to which this concept is not used directly to describe people but rather to do something like compliment (Rorty 1979; 1995), tag (Craig 1990), or keep the books on the ways in which various people track, process, and transmit information (Brandom 2000: ch. 3, Rosenberg 2002: ch. 5).

In an earlier paper (Chrisman 2007), I provisionally argued for a version of epistemic expressivism not on the basis of these parallels but on the basis of the fact that it addresses two pervasive objections to epistemic contextualism while retaining the advantages of that view. I want to leave this argument to the side here, except to say that what resulted was a view that encourages us to ask not about the nature of epistemic facts and values but instead about the nature of epistemic evaluations. It did so with the typical hope of expressivist views to answer the latter question without ever needing to countenance some *sui generis* kind of thing – epistemic facts or values – to answer the former question. The idea is the same as before: epistemic evaluations are not seen as representations of how the world is; they are seen as expressions of our attitudes towards the way we represent the world to be. In particular, they involve expressions of our distinctively epistemic attitudes towards things like believers, beliefs, and types of evidence and justification.

Just like before, from the perspective of truth-conditional semantics, the resulting semantics for epistemic claims is apt to seem quite radical. However, epistemic expressivists seem to have at least some support from ideationalist or psychologistic approaches to meaning on their side (as well, of course, ethical expressivists) in arguing that epistemic sentences could mean what they do in virtue of the state of mind that they express, rather than in virtue of what they represent.

The epistemic parallel to the toy version of ethical expressivism above would hold

⁹ Some expressivists such as Ayer (1936), Stevenson (1937), and Gibbard (1990) have assumed that this commits them to denying the truth-aptness of ethical claims. However, more recent expressivists such as Blackburn (1998), Timmons (1999), Gibbard (2003) and Ridge (2009) have sought to avoid this commitment in various ways.

¹⁰ One exception to this is ethical expressivists pursuing the "quasi-realist" program of reconstructing the language of realism within an antirealist expressivism. Thus, both Blackburn (1996: 87; 1998: 318) and Gibbard (2003: 235) argue briefly for extending their ethical expressivism to cover epistemic claims. Outside of that, Field (1998) and Heller (1999) have sketched views of particular epistemic claims that might be read as expressivist views, although they have not been worked out in very much detail.

¹¹ Compare citations in note 2 above.

¹² See for instance Kvanvig (2003), Pritchard (2007), and the papers collected in Haddock, Millar, and Pritchard, eds. (forthcoming).

something similar about attributions of justification and knowledge. For example, it might hold that to say ‘S is justified in believing that p’ is to express one’s acceptance of a system of epistemic norms which one believes to positively value S’s belief that p. And, sloughing over a few details, we might extend this by claiming that to say ‘S knows that p’ is to express one’s acceptance of a system of epistemic norms, which one takes to positively value S’s true belief that p. In addition to retaining advantages of contextualism while overcoming two of its problems (if the argument of my previous paper is correct), this view would carry at least one of the same theoretical advantages as ethical expressivism.

This is the second advantage mentioned above: Epistemic expressivism would avoid queerness/placement problems for epistemic facts and concepts. Philosophers have been less concerned about the ontological status of epistemic facts and values than ethical facts and values; however, arguably, whatever naturalistic inclinations motivate concern about the queerness/placement of ethical facts or concepts, the same naturalistic inclinations should motivate the same concerns about epistemic facts and concepts, insofar as we’re conceiving of these as genuinely normative (whatever, exactly *that* means). If that is right, then epistemic expressivism would avoid an ontological/conceptual cost in its contention that epistemic concepts function like ethical concepts to make claims that express attitudes rather than describe or represent the world.

The other advantage typically claimed for ethical expressivism was its intuitive explanation of the practicality of ethical discourse, and it is perhaps because epistemic discourse isn’t seen as similarly practical that epistemologists have ignored that accounts of the meaning of epistemic claims that don’t turn on the truth-conditions of these claims, such as expressivist accounts. After all, claiming, e.g. that Plato didn’t know things about Google doesn’t seem to be as tightly connected to being motivated to act in a particular way as claiming, e.g. that charity is good. However, although I think it’s probably correct that epistemic discourse is not practical in exactly the same way that expressivists claim that ethical discourse is practical, I suspect that epistemic discourse is interestingly practical in another sense that encourages rethinking the pervasive presupposition in epistemology about how best to approach an account of the meaning of epistemic claims.

More on this in §4, but first I want to discuss another – to my mind, better and more general – reason for being wary of epistemic expressivism as an account of the meaning of epistemic claims.¹³ A brief and somewhat dogmatic discussion of the notorious Frege-Geach problem in meta-ethics will help to expose this reason.

3. EXPRESSIVISM’S DISCONTENTS

There are other views going under the label ‘expressivism’ on the meta-normative scene¹⁴, but one of the main challenges for views, which hold that normative claims mean what they do in virtue of the type of desire-like attitude that they express, is to explain, in a systematic way, the logic of normative concepts in all of their uses.

¹³ Kvanvig (2002: ch. 7), Cuneo (2008: ch 5-6), and Lynch (forthcoming) have argued that epistemic expressivism faces special problems not faced by ethical expressivism. These arguments are discussed and rejected in Carter and Chrisman (ms).

¹⁴ For example, Copp (2001) defends a view he calls ‘realist expressivism’, where the expressive function of ethical claims is seen as deriving from something other than their semantic content. See also the papers cited in note 3 above.

Normative concepts are clearly used not only atomically but also in logically complex contexts. We say things not only like ‘X is good’ but also like ‘If x is good, then y is good’; not only like ‘S’s belief that p is justified’ but also like ‘If S’s belief that p is justified, then S’s belief that q is not justified’. If expressivists say that the atomic claims mean what they do in virtue of expressing a desire-like attitude, what will they say about the logically complex claims? It seems that one could endorse these claims without expressing a desire-like attitude (towards x or y, or towards S’s belief that p or S’s belief that q). This is the Frege-Geach problem (Geach 1965).

In my view, ethical expressivism in its early forms suffered the Frege-Geach problem because, in effect, it treated the meaning of normative predicates as given completely by the expressive force of applying them in making an atomic normative claim. This is why Ayer suggested that atomic moral claims like ‘Stealing money is wrong’ have no factual meaning and express no proposition. But such a view is hopeless when it comes to explaining the semantic properties of normative predicates embedded in force-stripping contexts such as conditionals. There must be some stable content between ‘Stealing money is wrong’ and the antecedent of ‘If stealing money is wrong, then it will be good thieves are punished’, if we’re to view a modus ponens inference proceeding from these premises to be innocent of the fallacy of equivocation.

Interestingly, Ayer’s view of complex normative sentences of another type may, if suitably generalized, provide some traction against this problem. For the case of a complex normative sentence, he suggested that the sentence has descriptive content that is presented in a particular way. His example is the claim ‘Your stealing that money was wrong’, and he suggests that it expresses the proposition that you stole that money, but it presents this proposition in a special way that could in principle be replaced by a special form of exclamation marks indicating a speaker’s contempt for the action mentioned. What’s interesting is that this can help to capture some stable content across unembedded and embedded contexts for such complex claims. For Ayer can say that the same proposition is the content of the antecedent of ‘If your stealing that money was wrong, then you will be punished’.

A problem with this, of course, is that we need to explain the content of not only complex moral claims but also atomic moral claims. But many latter-day expressivists have sought to overcome this problem by also holding that atomic claims have *both* descriptive and evaluative aspects.¹⁵ It is highly controversial whether any of these enriched expressivist views work, but, rather than get into the trenches of that debate, I want to point out that, insofar as these views continue to operate within the general ideationalist approach to meaning, they can still be fit into Ayer’s heuristic at least this

¹⁵ In addition to the quote from Gibbard in note 6 above, compare Gibbard (1990: 93) and Ridge (2006) for the most explicit examples of this strategy. Both, however, echo Hare’s (1952: 18-20) distinction between the neustic and phrastic elements of meaning. One possible exception to this is Blackburn (1988/1993), who argues that logically complex sentences serve a particular commissive function; that is, they ‘tie one to a tree’ of conditional commitments. For example, the claim ‘If your stealing that money was wrong, then you will eventually be punished’ commits its author to the content of its consequent if she is committed to the content of its antecedent, and to the negation of its antecedent if she is committed to the negation of its consequent. I think approaches of this sort are quite promising; however, they seem to me to rest on a tacit abandonment of the ideationalist alternative to truth conditional semantics, from which expressivism seemed to receive some shelter from objections to its implicit semantic assumptions. For to hold that logically complex statements involving normative predicates mean what they do in virtue of ‘tying one to a tree’ of conditional commitments is to endorse something much more like a conceptual or inferential role semantics. If that’s right, then the resulting view may be consistent with the advantages traditionally claimed for expressivism (more on this below), but I think it can no longer be fairly thought of as a species of the view that ethical claims mean what they do in virtue of the mental states that they express.

far: normative sentences are still viewed as having a particular descriptive content that is presented in a special normative way that could in principle be replaced with a special set of exclamation marks expressing the relevant desire-like attitude.

For example, we can already see how this is the case with our toy version of ethical expressivism above. The idea was that a sentence of the form

‘X is good’

expresses the acceptance of a set of norms N, which one takes to encourage promoting x. Here the descriptive content is that

N encourages promoting x

and it is presented, so to speak, in the mode of one who accepts N. Thus, on the view, it seems that we could in principle replace the claim ‘X is good’ with ‘N encourages promoting x!!’, where the ‘!!’ indicates that the speaker accepts N. When the normative predicate is embedded in a force-stripping context, the point remains the same. For instance, the sentence

‘If x is good, then y is good’

might be thought to have the descriptive content that

if N encourages promoting x, then N encourages promoting y

which is presented in the mode of one who accepts N. Thus, it could, in principle be replaced with ‘If N encourages promoting x, then N encourages promoting y!!’, where the ‘!!’ indicates that the speaker accepts N.

No mature version of expressivism deploys Ayer’s heuristic tactic so baldly, but I think most can be fit into the general model.¹⁶ Because of this, however, I worry that they trade on an impoverished view of what a normative concept could be. To appreciate this, let me contrast two conceptions of concepts.¹⁷

There is a strong tradition in cognitive science of viewing concepts, especially those often deployed in observation as asymmetrically dependent labels or classificatory devices. On this view, concept possession is a discriminatory ability: to possess the concept C is to be reliably disposed to respond differentially to things in a particular class, the C-things. For instance, the concept ‘red’ is, on this view, a classificatory device for the red things; to possess this concept is to be reliably disposed to respond differentially to red things (under normal circumstances and with the right sort of prompting) with the word ‘red’ (or some suitable analogue), i.e. to be able to classify things as red or non-red.¹⁸

Another tradition dominates the philosophy of logic, where concepts are seen as logically articulable functions or nodes of inferential (evidential and/or justificatory) relations. On this view, concept possession is a rational ability: to possess the concept C

¹⁶ One possible exception is the version of expressivism explored in Schroeder (2008), which seeks to reconstruct the compositionality and semantic properties of normative statements in a general expressivist framework. However, in order to capture the semantic properties, he ends up forcing the expressivist to reconstruct notions of validity and entailment for all statements (both normative and nonnormative) in a global expressivist account of language – that is, all indicative sentences are thought to mean what they do in virtue of the desire-like attitude of ‘being-for’ that they express. Even if this view escapes the criticism I go on to levy in the text above, it is not a very plausible view. As Schroeder himself concludes, it is ‘an extremely unpromising hypothesis about the workings of natural language’ (2008: 179).

¹⁷ This is a simplification of a distinction deployed in Brandom (2009) drawing on Sellars (1956).

¹⁸ For philosophical expositions of theories of concepts which are at least inspired by this general tradition, see Dretske (1981) and Fodor (1990), though both Dretske and Fodor pursue more complex models, anticipating many of the most natural objections to the crude summation of the tradition presented in the text above.

is to be able to draw the right sorts of inferences to and from applications of the concept C. For instance, the concept of conjunction is, on this view, the node of the inferential relations between judgments deploying this and other concepts; to possess this concept is to have the ability to make (some substantial portion) of these inferences—e.g. from ‘P and Q’ to ‘Q’ and from ‘P’ and ‘Q’ to ‘P and Q’.¹⁹

Which conception of concepts is correct? I think both conceptions capture something that is important about at least some concepts, so I’d say neither is correct if it is meant to capture all that there is to concepts and concept possession. However, when it comes to the nature of *human* concepts and especially our normative concepts, I think the inferentialist conception has some claim to being more illuminating and fundamental. For one thing, it doesn’t have to *deny* that concepts sometimes serve as classificatory devices, it just claims that there’s more to human concepts. For another, when concepts have a richer logic than is captured by their possible use as classificatory devices, this view can explain this logic. This seems to be especially apt when it comes to logical concepts such as ‘necessary’, but it may extend to other concepts such as normative concepts too. Even those concepts, which are unquestionably observational, might be thought to embody inferential nodes. For example, it’s plausible to think ‘x is red’ implies ‘x is colored’ and ‘x is crimson’ implies ‘x is red’. And this could be relevant for determining their classificatory function.

I bring up this distinction between the two views of concepts because I suspect that expressivists are, to put things somewhat crudely, tacitly conceiving of normative concepts as classificatory devices rather than as nodes of inferential relations. It’s not that they think that a normative concept such as ‘good’ is analogous to an observational concept such as ‘red’ in bearing an asymmetric relationship to all and only the good things. That would be a form of realist cognitivism. Rather, it seems to me that they treat ethical concepts such as ‘good’ as partially second-order devices for classifying the things towards which we (the concept-users) have a certain desire-like attitude in virtue of certain descriptive properties of the things.

In other words, although, on the expressivist view, ‘x is good’ clearly doesn’t describe x as something the speaker approves of, it does classify x as one of the things to which the speaker has a specific desire-like attitude of approval. The same is true even in logically complex contexts such as a conditional. For example, on the expressivist view of the sentence ‘If x is good, then y is good’, neither x nor y is classified as unconditionally good. But they are classified as standing in a certain relation with respect to the things the speaker approves of – basically, if x is in the class of things the speaker approves of then so is y. This is why we can reconstruct the expressivist’s view of normative sentences as one according to which they have descriptive content that is presented in a special way, which we could in principle replace with a certain kind of tag or label, e.g. the special exclamation marks from Ayer.

On this way of thinking of things, to possess a normative concept is to be disposed to respond differentially to the things towards which one has a certain desire-like attitude (under normal circumstances and with the right sort of prompting) with the word ‘good’ (or some suitable analogue). This is why the expressivist has such an easy time explaining the intuitive practicality of normative discourse, as explained above. However, just as I think the general conception of concepts as classificatory device will have a hard time explaining, in general, the logic of human concepts, I think the view of ethical concepts which seems to be tacit in expressivism, is likewise bound to have a hard time explaining

¹⁹ For philosophical expositions of theories of concepts which are inspired by this general tradition, see Sellars (1956), Dummett (1974), Rosenberg (1974), and Brandom (1994; 2000; 2009).

the logic of ethical concepts.

There may be a structurally adequate expressivist response to the Frege-Geach problem. That is to say that there may be a version of the expressivist idea that normative claims express desire-like attitudes, which is capable of assigning different mental states to different sentences as a way of giving their meaning at whatever degree of logical complexity. However, as long as the aspect of normative sentences that makes them normative is conceived of as a second-order desire-like classificatory device, I think expressivists will be stuck with an impoverished view of normative concepts. This is not to say that it is wrong to view normative concepts as (partially) second-order classificatory devices, for that idea may very well capture an important aspect of these concepts. However, especially in the normative case, it seems to me that this cannot be the most illuminating or fundamental story, since it doesn't explain the inferential (evidential or justificatory) relations in which normative concepts stand.

This is, I think a better reason for being wary of expressivism, and it may seem to be even more pressing in the case of epistemic expressivism than in the case of ethical expressivism. For perhaps philosophy is still influenced by positivist undercurrents enough to take seriously the idea that ethical concepts are somehow attitudinal classificatory devices rather than inferential nodes. However, positivism's influence hasn't extended far enough to make the idea that epistemic concepts are attitudinal classificatory devices rather than inferential nodes equally acceptable. Epistemic concepts clearly have a robust logic, and it seems to be one entirely parallel to other nonnormative concepts. This is something that I think any view of normative concepts, which turns implicitly on the cognitive science conception of concepts rather than the philosophy of logic conception of concepts, is bound to fail to explain.

4. INFERENCEALISM AS A NEW META-ETHICAL AND META-EPISTEMOLOGICAL VIEW

Summing up before moving on: Meta-epistemology has largely presupposed that debates about the meaning of epistemic claims will be settled by determining the truth-conditions of these claims. I've claimed that this obscures meta-normative views like expressivism which explain the meaning of the target claims in terms other than truth-conditions. However, in both meta-ethics and meta-epistemology, although expressivism claims certain theoretical advantages, it also seems to tacitly assume a view of normative concepts that will not be taken to be the most fundamental story in the epistemic case and, anyway, is impoverished enough to wonder whether it should be taken as the most fundamental account of any concepts at all.

In light of the distinction between the two views of concepts, I now want to suggest a way to retain the advantages of expressivism without incurring the problem I've raised for it. The strategy will be to pick up on the idea of concepts as inferential nodes and to sketch the rudiments of a new meta-normative view with this at its heart. As a first pass, this view can be seen as involving three related elements.

First, it says that we should think of normative concepts not (merely) as classificatory devices (even at the second-order) but as nodes of inferential relations. Which inferential relations? It is difficult to say in general and exhaustively, but these will comprise all of the standard logical relations (like *modus ponens* and *modus tollens* inferences) and the conceptual connections between these concepts and other concepts

(e.g. ought implies can, knowledge implies truth, etc.).²⁰

Second, it says that we should recognize the distinctively practical kinds of inferential relations in which some normative concepts are caught up. For example, while ‘S ought all things considered to do A’ may have the theoretical implication that ‘S can do A’, from the first-person, ‘I ought all things considered to do A’ may have the practical implication ‘I shall do A’ (where this is not the expression of a prediction but rather the expression of an intention). Likewise, while ‘S knows that p’ may have the theoretical implication ‘S believes that p’; from the first person, ‘I know that p’ may have the practical implication ‘I shall act as if p is true’ and/or ‘I shall stop inquiring as to whether p’ (where this is not the expression of a prediction but rather the expression of an intention). In effect, here we get a distinction between two different sorts of inferential roles: the theoretical inferential roles that normative concepts may share with descriptive concepts and the practical inferential roles that mark them out as different.

Third, it says that we should retain, if only provisionally, the popular idea that intention and belief have different directions of fit with the world, such that, if two people believe incompatible propositions to be true, their disagreement presupposes some fact of the matter about who is right, whereas, if two people intend to make different propositions true, their disagreement doesn’t by itself presuppose some fact of the matter about who is right.

Now, if we hold that normative predicates and the sentences in which they figure mean what they do in virtue of their inferential role, so construed, I think we get an alternative to both truth-conditionalism and expressivism, as accounts of the meaning of normative claims. We might call this view *normative inferentialism* (with *ethical inferentialism* and *epistemic inferentialism* as two species).²¹ Its attractiveness depends on its ability to retain the advantages of expressivism while enriching the conception of normative concepts in a way that avoids expressivism’s problems. I’ve just said how it enriches the conception of normative concepts by moving beyond classificatory devices to inferential nodes. I’ll illustrate separately how it might retain the advantages of ethical expressivism and then turn to epistemic expressivism.

In the meta-ethical debate, expressivists claim an easy explanation of the intuitive practicality of ethical discourse. However, above, I distinguished two senses in which ethical discourse is distinctively practical. There is, an issue in, so to speak, the order of causes: those who make certain sorts of ethical claims tend to exhibit correlated motivations. By contrast, there is an issue in the order of reasons: one can, it seems, make ethical claims to justify actions. I think ethical inferentialism is consistent with many different explanations of the first sense in which ethical discourse is practical. Indeed, it is consistent with the idea that ethical claims express desire-like attitudes that serve, in the order of causes, to motivate the expected action (although, it won’t put this idea to the same work in the theory of meaning that the expressivist does). However, ethical inferentialism offers a novel explanation of the second sense in which ethical discourse is distinctively practical. By treating ethical predicates as standing in inferential relations with the predicates that serve in the expression of intentions, we can explain more than just why ethical claim-making tends to cause certain motivation, we can

²⁰ In this way, the account bears some similarity to Wedgwood’s (2007: ch. 4-5) conceptual-role semantics for normative claims. However, I differ from Wedgwood in thinking that conceiving of meaning in terms of inferential-roles opens up room for antirealist views in normative ontology rather than forces one to adopt a very strong form of normative realism.

²¹ I leave it open here whether, like truth-conditionalism, inferentialism should be seen as a global view of meaning. If it is, normative inferentialism is an application, which draws a contrast between the normative and the non-normative in terms of type of inferential role.

explain why it can be used in an attempt to justify it. The idea is that it is by standing in inferential relations to intentions (to action) that ethical claims are of the right inferential category to be made in response to demands for justification of an action. For example, if you give some money to charity, and then someone asks you, 'Why did you do that?', you can typically say, 'Because giving to charity is a good thing to do.' That is, you can answer, in the order of reasons, by making an ethical claim. Ethical inferentialism explains why this is so.²²

The other main advantage of ethical expressivism was its claim to avoid queerness/placement worries generated by the practicality of ethical concepts. Here, it's difficult to assess the relative merit of inferentialism because it's difficult to say when the inferentialist semantic idea, viz., that claims mean what they do in virtue of their inferential roles, commits one ontologically. Some inferentialists (e.g. Dummett) seem attracted to antirealism across the board, in which case ethical inferentialism would obviously be just as ontologically parsimonious as ethical expressivism, but only on pain of deconstructing the whole issue of ontology. Other inferentialists (e.g. Sellars, Harman) seem willing to reconstitute the notion of representing the world within the inferentialist framework. However exactly they do this, as long as there is room for viewing some claims as not completely representational because of their inferential relations to expressions of intentions, ethical inferentialism can argue that ethical claims fit into this latter category and thereby retain the advantage of ontological parsimony claimed by ethical expressivists. The idea would be to think of some concepts as bearing no direct inferential connection to intentions, and so not having a distinctively practical inferential role. These concepts could be viewed as the descriptive or representational concepts, whose deployment in basic assertions commits the speaker to the existence of the corresponding fact.²³ For example, 'x is red' could then be viewed as committing one to the fact that x is red. By contrast, other concepts could then be thought of as lacking this commitment precisely because of their distinctively practical inferential role (and the corresponding difference in 'directions of fit'). Assuming that ethical concepts belong in the latter category, ethical inferentialism would have a parallel way to avoid queerness/placement worries to the ethical expressivist.²⁴

The ostensible advantages of an expressivist approach have not seemed as weighty in the meta-epistemological debate about the meaning of epistemic claims as they may seem in the meta-ethical debate. However, in the case of queerness/placement worries, I think this may have resulted from mere oversight. If we conceive of them as normative, epistemic facts and values seem to me to be just as difficult to locate in the natural world as ethical facts and values. For one thing, epistemic concepts seem to resist analysis in nonnormative terms. For another, epistemic norms seem to be categorical and not hypothetical, in the sense that their normative force doesn't depend on the contingent and personal aims or desires of the agents to whom they apply.²⁵ In any case, to the

²² You might also be able to answer in the order of causes by saying 'Because I approve of giving to charity.' This exhibits the ambiguity of the question 'Why did you do x?' Ethical inferentialism as I am conceiving of it is consistent with this response in the order of causes, but it denies that such a response is adequate in the order of reasons.

²³ At least, this could be viewed as a necessary condition for such commitment. Things may be more complicated depending on one's view of ontological commitment.

²⁴ What about views which see ethical claims as having both practical and theoretical inferential connections? On the present account, as long as a claim like 'x is good' commits one to 'I shall promote x', it has a practical inferential role and so does not commit one ontologically to the fact that x is good. However, it might imply, e.g. that x is approved of by many people, and this then could commit one to the fact that x is approved of by many people.

²⁵ Compare Kelly (2003) and my discussion of Kornblith (2001) in Chrisman (2008: 354-8).

extent that *ethical* inferentialism retains resources for claiming ontological parsimony over truth-conditionalist views of the meaning of ethical claims, *epistemic* inferentialism will do so as well. So, if epistemic facts and values are queer enough to make ontological commitment to them a theoretical liability, epistemic inferentialism will retain some of the advantage of epistemic expressivism.

As I mentioned above, the expressivist's easy explanation of the motivational force of ethical discourse, does not seem to carry over to epistemic discourse. We can, of course, make some pretty reliable predictions about what someone who affirms certain epistemic claims will do. For example, there may be a tight connection between claiming 'I know that p' and being motivated to act as if p is true. However, it seems that this has nothing specifically to do with the claim's being an epistemic claim. We can just as well make such predictions from someone who asserts that p.²⁶

Nonetheless, it seems to me that epistemic discourse is practical in a different sense. As with ethical claims, I think epistemic claims can be made in the attempt to justify certain actions. Which epistemic claims can be used to justify which actions is a matter of considerable debate. Some epistemologists think knowing that p justifies asserting that p or acting as if p is true, while others think being justified in believing p is enough to justify these actions.²⁷ Whatever exactly the correct account is, if some epistemic notions serve as the norm of some kinds of actions, then epistemic discourse exhibits a sort of practicality in the order of reasons that demands explanation. Epistemic inferentialism provides this explanation by construing epistemic predicates as inferentially related to the predicates by which we express intentions.

Thus, it seems to me that normative inferentialism – with its ethical and epistemic species – will be at least as well off as normative expressivism vis-à-vis the alleged advantages of the latter view. And, since normative inferentialism is based on a view of concepts, including especially ethical and epistemic concepts, as inferential nodes rather than classificatory devices, it will overcome one of the core drawbacks of normative expressivism: its impoverished conception of normative concepts.

5. TOWARDS A VIABLE SOCIAL EPISTEMOLOGY IN PRAGMATISM

In many ways, the sort of epistemic inferentialism for which I have been arguing here is inspired by a sort of pragmatism about words and concepts that encourages us to look for their social-practical roles rather than referents in order to understand their meaning. However, it's a pragmatism which makes a distinction between normative concepts like 'knows' or 'good' and descriptive concepts like 'water' or 'heat'. The difference will be not in which concepts have inferential connections but in the inferential connections between these concepts and action. In my view, normative concepts embed the concept 'ought' and this we should understand in terms of its inferential connections to intentions to action. I think this allows us to view normative concepts as playing some essentially nondescriptive social-practical role.

Because other pragmatists sometimes say things sounding roughly like this, I'd like

²⁶ Of course, on some accounts of assertion, asserting that p commits one to 'I know that p'; however, even if that's right, it's still unclear that the implicit knowledge claim is providing the link to motivation rather than the explicit claim that p.

²⁷ Consider the debate over whether knowledge is a 'norm of action' among, for instance, Pollack and Cruz (1999: ch. 5), Williamson (2005), Stanley (2005), Hawthorne and Stanley (2008), Neta (2008), Brown (2008).

to conclude by distinguishing the pragmatist social epistemology an epistemic inferentialist should allay himself with from some other similar ideas.

As I mentioned above, pragmatists often challenge the standard approach to doing meta-epistemology. For example, Rorty (1979, 1995) suggests that epistemic concepts are primarily terms of praise. This allows him to cast the endless debate about the nature of knowledge as stemming from a failure to recognize that ‘knows’, like ‘good’, is not a word that picks out something with a nature. In a somewhat different vein, Craig (1990) proposes to approach the concept of knowledge by investigating what role the introduction of such a concept would have served in human coordination in community, much like Hobbes tried to understand the concept of justice by investigating what role the introduction of such a concept would have served in human coordination in community. Thus, Craig’s account of ‘knows’ construes it as (initially) a way to tag reliable informants; what it takes for someone to be deserving of such a tag will largely be down to the varying specific information tracking and transmitting needs of different groups of people and not some underlying ‘nature of knowledge’.²⁸

Both of these approaches to understanding epistemic concepts have met significant resistance, which is I think due to the vague sense that they undermine the objectivity of epistemic concepts and discourse. If epistemic terms are used merely to praise people’s beliefs in a special epistemic way, and different people find different sorts of things to be worthy of epistemic praise, then there will be no objective answer to first-order epistemic questions. Similarly, if epistemic terms are used to tag reliable informants, but there is no specific measure of reliability one must achieve to deserve such a tag, then there will be no objective standard for our epistemic claims.

Despite this worry, I think it’d be wrong to doubt that epistemic concepts are used to praise and tag; it’s just that they seem to do or be something more. If the view of epistemic concepts as devices of praise or tagging encourages something like the expressivist semantics for epistemic claims that I discussed earlier, I think the inferentialist semantics that I have proposed to view as expressivism’s successor (both in meta-ethics and meta-epistemology) offers a more nuanced view of the essentially social function of epistemic concepts and, thereby, a way to view them as doing more than merely praising and tagging.

The idea would be to treat epistemic concepts as nodes of inferential relations, which are ossified and perpetuated by the social-linguistic practice of users of these concepts in human linguistic community as they faced certain needs. Thus, there will be good pragmatic reasons to explain why it is better for us to use concepts embodying these inferential roles rather than some other concepts embodying some other inferential roles. However, once we are using the concepts we have good reason to use, since they embody specific inferential roles, there is an objective answer to questions about when they can be correctly applied and what follows from such an application. (Compare: an inferentialist about logical concepts usually holds that there is an objective answer about the correct application of, and what follows from a conjunction or negation.) This is why I have taken my inspiration from other pragmatists about knowledge, like Brandom (2000: ch. 3) and Rosenberg (2002: ch. 5-6) who embed their pragmatism about epistemic concepts in a broader pragmatism about concept possession and use more generally that underwrites an inferentialist treatment of these concepts.

The objectivity for which this move makes room is not the objectivity of some sort

²⁸ These ideas are related to Austin’s suggestion that ‘When I say ‘I know’, I give others my word’(1979: 99). However, I take it Rorty and Craig want their ideas to extend to third-personal ascriptions of knowledge as well.

of Realism about the relevant kinds of predicates and concepts that they express (e.g. moral, epistemic, logical, etc.). Rather, it's the kind of objectivity sought by Kant's discussion of pretty much all of the concepts that interested him. His idea is that we can establish our unconditional right to use the concepts without needing to establish access to properties in the world that these concepts purport to represent. The value I see in treating this as stemming from the social function of epistemic discourse is that it encourages investigation into why it is useful to use epistemic concepts the way that we do. And this is the sort of investigation that single-minded focus on the truth conditions of epistemic claims (invariantist or contextualist) or even on the expressive role of epistemic claims (representationalist or expressivist) is likely to hinder rather than help.²⁹

²⁹ I'd like to thank an audience at the University of St. Andrews as well as J. Adam Carter, Georgi Gardiner, Graham Hubbs, Alan Millar, and Michael Ridge for helpful feedback on earlier versions of this material.

WORKS CITED

- Austin, J. L. (1979) *Philosophical Papers*, Oxford: Oxford University Press.
- Ayer, A. J. (1946) *Language, Truth and Logic*, London: V. Gollancz Ltd.
- Bar-On, Dorit & Matthew Chrisman (2009) 'Ethical Neo-Expressivism', in Shafer-Landau, Russ (ed.), *Oxford Studies in Metaethics, Volume 4*, Oxford; New York: Oxford University Press: 133-165.
- Blackburn, Simon (1984) *Spreading the Word: Groundings in the Philosophy of Language*, Oxford; New York: Clarendon Press; Oxford University Press.
- Blackburn, Simon (1988/1993) 'Attitudes and Contents', reprinted in *Essays in Quasi-Realism*, New York: Oxford University Press.
- Blackburn, Simon (1993) *Essays in Quasi-Realism*, Oxford; New York: Oxford University Press.
- Blackburn, Simon (1996) 'Securing the Nots: Moral Epistemology for the Quasi-Realist', in Sinnott-Armstrong, Walter & Mark Timmons (eds.), *Moral Knowledge: New Readings in Moral Epistemology*, Oxford; New York: Oxford University Press: PAGES.
- Blackburn, Simon (1998) *Ruling Passions: A Theory of Practical Reasoning*, Oxford; New York: Clarendon Press; Oxford University Press.
- Blackburn, Simon (2006) 'The Semantics of Non-Factualism', in Devitt, Michael & Richard Hanley (eds.), *The Blackwell Guide to the Philosophy of Language*, Malden, MA; Oxford: Blackwell: 244-252.
- Brandom, Robert (1994) *Making It Explicit: Reasoning, Representing, and Discursive Commitment*, Cambridge, MA: Harvard University Press.
- Brandom, Robert (2000) *Articulating Reasons: An Introduction to Inferentialism*, Cambridge, MA: Harvard University Press.
- Brandom, Robert (2009) *Reason in Philosophy: Animating Ideas*, Cambridge, MA: Harvard University Press.
- Brown, Jessica (2008) "'Subject-Sensitive Invariantism" and the Knowledge Norm of Practical Reasoning', *Noûs* 42: 167-189.
- Carter, J. Adam and Matthew Chrisman (ms) 'Is Epistemic Expressivism Incompatible with Inquiry?'
- Chrisman, Matthew (2008) 'From Epistemic Contextualism to Epistemic Expressivism', *Philosophical Studies* 135: 225-254.
- Chrisman, Matthew (2008) 'Ought to Believe', *Journal of Philosophy* CV: 346-370.
- Chrisman, Matthew (2009) 'Expressivism, Truth, and (Self-)Knowledge' *Philosophers' Imprint*, 9(3): 1-26.
- Chrisman, Matthew (ms) 'Metaethics and the Philosophy of Language'.
- Cohen, Stewart (1988) 'How to Be a Fallibilist', *Philosophical Perspectives* 2: 91-123.
- Cohen, Stewart (1999) 'Contextualism, Skepticism, and the Structure of Reasons', *Noûs* 33: 57-89.
- Copp, David (2001) 'Realist-Expressivism: A Neglected Option for Moral Realism', *Social Philosophy and Policy* 18: 1-43.
- Craig, Edward (1990) *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Oxford: Clarendon Press.
- Davidson, Donald (1980) *Actions and Events*, Oxford: Oxford University Press.
- Davis, Wayne (2003) *Meaning, Expression, and Thought*, Cambridge: Cambridge University Press.
- DeRose, Keith (1992) 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological*

Research 52: 913-29.

- DeRose, Keith (1995) 'Solving the Skeptical Problem', *The Philosophical Review* 104: 1-52.
- DeRose, Keith (2009) *The Case for Contextualism: Knowledge, Skepticism and Context, Volume I*, Oxford; New York: Oxford University Press.
- Dretske, Fred (1981) *Knowledge and the Flow of Information*, Cambridge, MA: Harvard University Press.
- Dummett, Michael (1974) *Frege: Philosophy of Language*, Cambridge, MA: Harvard University Press.
- Field, Hartry (1998) 'Epistemological Nonfactualism and the a Prioricity of Logic', *Philosophical Studies* 92: 1-24.
- Fodor, Jerry (1990) *A Theory of Content*, Bradford, MA: MIT Press.
- Geach, P. T. (1965) 'Assertion', *The Philosophical Review* 74: 449-465.
- Gibbard, Allan (1990) *Wise Choices, Apt Feelings: A Theory of Normative Judgment*, Cambridge, MA: Harvard University Press.
- Gibbard, Allan (2003) *Thinking How to Live*, Cambridge, Mass.: Harvard University Press.
- Grice, H. P. (1957) 'Meaning', *The Philosophical Review* 66: 377-388.
- Haddock, Adrian, Alan Millar, & Duncan Pritchard, eds. (forthcoming) *Epistemic Value*, Oxford; New York: Oxford University Press.
- Hare, R. M. (1952) *The Language of Morals*, Oxford: Clarendon Press; Oxford University Press.
- Hawthorne, John (2004) *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Hawthorne, John & Jason Stanley (2008) 'Knowledge and Action', *Journal of Philosophy* 105: 571-590.
- Heller, Mark (1999) 'The Proper Role for Contextualism in an Anti-Lick Epistemology', *Philosophical Perspectives* 13: 115-129.
- Kelly, Thomas (2003) 'Epistemic Rationality as Instrumental Rationality: A Critique', *Philosophy and Phenomenological Research* 66: 612-640.
- Kornblith, Hilary (2001) 'Epistemic Obligation and the Possibility of Internalism', in, A. Fairweather and L. Zagzebski (eds.), *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, New York; Oxford: Oxford University Press.
- Kvanvig, Jonathan (2003) *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.
- Lewis, David (1996) 'Elusive Knowledge', *Australasian Journal of Philosophy* 74: 549-67.
- Mackie, J. L. (1977) *Ethics: Inventing Right and Wrong*, London: Penguin.
- Neta, Ram (2003) 'Contextualism and the Problem of the External World', *Philosophy and Phenomenological Research* 66: 1-31.
- Neta, Ram (2008) 'Anti-Intellectualism and the Knowledge-Action Principle', *Philosophy and Phenomenological Research* 75: 180-187.
- Pollock, John & Joseph Cruz (1999) *Contemporary Theories of Knowledge*, Lanham, MD, Rowman & Littlefield Publishers.
- Pritchard, Duncan (2007) 'Recent Work on Epistemic Value', *American Philosophical Quarterly* 44: 85-110.
- Ridge, Michael (2006) 'Ecumenical Expressivism: Finessing Frege', *Ethics* 16: 302-336.
- Ridge, Michael (forthcoming) 'Truth in Ecumenical Expressivism' in Sobel, ed. *Reasons for Action*, Cambridge University Press: Cambridge.
- Rorty, Richard (1979) *Philosophy and the Mirror of Nature*, Princeton, NJ: Princeton University Press.

- Rorty, Richard (1995) 'Is Truth a Goal of Enquiry? Davidson Vs. Wright', *Philosophical Quarterly* 45: 281-300.
- Rosenberg, Jay (1974) *Linguistic Representation*, Dordrecht: D. Reidel Publishing Co.
- Rosenberg, Jay F. (2002) *Thinking About Knowing*, New York, Oxford: Clarendon Press; Oxford University Press.
- Rysiew, Patrick (2001) 'The Context-Sensitivity of Knowledge Attributions', *Noûs* 35: 477-514.
- Schiffer, Stephen (1972) *Meaning*, Oxford: Oxford University Press.
- Schroeder, Mark (2008) *Being For*, Oxford; New York: Oxford University Press.
- Searle, John R. (1962) 'Meaning and Speech Acts', *The Philosophical Review* 71: 423-432.
- Sellars, Wilfrid (1956) 'Empiricism and the Philosophy of Mind', *Minnesota Studies in the Philosophy of Science*, volume 1, Feigl & Scriven (eds.), Minneapolis, MN: University of Minnesota Press: 253-329.
- Smith, Michael (1987) 'The Humean Theory of Motivation', *Mind* 96: 36-61.
- Smith, Michael (1995) *The Moral Problem*, Oxford: Blackwell.
- Sosa, Ernest (1980) 'The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge', *Midwest Studies in Philosophy* 5: 3-25.
- Sosa, Ernest (2007) *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford; New York: Oxford University Press.
- Stanley, Jason (2005) *Knowledge and Practical Interests*, Oxford; New York: Clarendon Press; Oxford University Press.
- Stevenson, Charles L. (1937) 'Emotive Meaning of Moral Terms', *Mind* 46: 14-31.
- Timmons, Mark (1999) *Morality without Foundations: A Defense of Ethical Contextualism*, New York: Oxford University Press.
- Unger, Peter (1984) *Philosophical Relativity*, Minneapolis: University of Minnesota Press.
- Wedgwood, Ralph (2007) *The Nature of Normativity*, Oxford; New York: Oxford University Press.
- Williams, Michael (1992) *Unnatural Doubts*, Princeton, NJ: Princeton University Press.
- Williams, Michael (2001) *Problems of Knowledge*, Oxford; New York: Oxford University Press.
- Williamson, Timothy (2000) *Knowledge and Its Limits*, Oxford; New York: Oxford University Press.
- Williamson, Timothy (2005) 'Knowledge, Context, and the Agent's Point of View', in Preyer, Gerhard, and Georg Peter (ed.), *Contextualism in Philosophy*, New York; Oxford: Oxford University Press.
- Zagzebski, Linda (1996) *Virtues of the Mind*, Cambridge: Cambridge University Press.