THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Constructivism, Expressivism and Ethical Knowledge

Citation for published version:

Chrisman, M 2010, 'Constructivism, Expressivism and Ethical Knowledge', *International Journal of Philosophical Studies*, vol. 18, no. 3, pp. 331-353. <https://doi.org/10.1080/09672559.2010.492119>

Digital Object Identifier (DOI):

[10.1080/09672559.2010.492119](https://doi.org/10.1080/09672559.2010.492119)

Link:

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

Document Version:

Peer reviewed version

Published In:

International Journal of Philosophical Studies

Publisher Rights Statement:

© Matthew Chrisman, 2010. Chrisman, M. (2010). Constructivism, Expressivism and Ethical Knowledge. *International Journal of Philosophical Studies*, 18(3), 331-353doi: 10.1080/09672559.2010.492119

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.

Constructivism, Expressivism and Ethical Knowledge

Matthew Chrisman

Forthcoming in a special volume of *International Journal of Philosophical Studies*, James O'Shea, ed. in conjunction with Naturalism, Normativity, and the Space of Reasons Conference at University College Dublin

ABSTRACT: In the contemporary metaethical debate, expressivist (Blackburn, Gibbard) and constructivist (Korsgaard, Street) views can be viewed as inspired by irrealist ideas from Hume and Kant respectively. One realist response to these contemporary irrealist views is to argue that they are inconsistent with obvious surface-level appearances of ordinary ethical thought and discourse, especially the fact that we talk and act as if there is ethical *knowledge*. In this paper, I explore some constructivist and expressivist options for responding to this objection. My conclusion is that, although both constructivists and expressivists can capture other surface-level features of ethical thought and discourse, the possibility of ethical knowledge causes special problems for these versions of irrealism. I end with some comments about where I think irrealists should begin to look for a response to these special problems, which points, somewhat surprisingly, towards an alternative inferentialist form of irrealism about epistemic and ethical thought and discourse, which is inspired by Sellars.

Introduction

One of the original motivations for irrealism in metaethics was skepticism about ethical properties. The argument is roughly as follows: According to a dominant empiricism tracing back at least to Hume and Kant, discovering the matter-of-factual properties of things involves those properties bearing some explanatory connection to possible objects of experience. Yet, if ethical properties existed, they wouldn't be like that. They wouldn't be like that because ethical values appear to bear some special 'internal' connection to action. And no empirically discoverable property could do that.

Each step of this argument is controvertible. However, assuming that we accept the argument, this skepticism about the putative ethical properties of things can seem to force us very quickly to endorse a radical error theory or a pernicious relativism in our metaethics. Irrealist views, such as expressivism and constructivism, offer a way to avoid this empiricist-driven slide into error-theory and relativism. They do so by encouraging a shift in initial focus from the putative objects of ethical thought and discourse to the subjects engaged in the practices that constitute such thought and discourse.¹ Then, here, they suggest we find not *description* of the natural world but rather something like the *expression* of moral attitudes or the *construction* of solutions to practical problems.

In this regard, Hume and Kant may each be read as already suggesting such accounts. Hume suggests that ethical thought and discourse is not mainly about putatively ethical matters of fact. (How could we ever have impressions of such things?) Instead he suggests ethics is mainly about bringing our sentiments into line with one another in pursuit of

¹ Compare Price (2004: 73) who makes a distinction between "object naturalism" and "subject naturalism".

communal cooperation and fellow feeling. Kant also suggests that ethical thought and discourse is not mainly about putatively ethical matters of fact. (How could such things be subject to the causal laws governing the world of things as they appear to us?) Instead he suggests that ethics is about how to align our wills with the logical laws we give ourselves in reasoning practically.

Each of these views offer an explanation of the apparently internal connection between ethics and action which turn on denying that ethical thought and discourse is best understood as describing some ordinary matter of fact. However, the explanations rest on markedly different understandings of the ‘internal’ connection. Followers of Hume are typically impressed by the apparently *regular* connection between one’s ethical judgments and one’s motivation to act in their stead. Roughly, if someone judges that she ought to ϕ , then *ceteris paribus* it’d be irregular if she weren’t at least somewhat motivated to ϕ . And because Humeans think that a motivation to act in one way rather than another must involve more than mere matter-of-factual beliefs, they suggest that ethical thought and discourse is more a matter of whatever else must be involved. In their view, that is something essentially desire-like, such as sentiments or attitudes. In contrast, followers of Kant are impressed by the apparently *rational* connection between ethical judgments and actions. Roughly, if someone judges that she ought to ϕ , then *ceteris paribus* it seems like she is practically irrational if she doesn’t ϕ . And because Kantians think that what one is rational to do always depends on more than the matter-of-factual beliefs that are within the scope of theoretical reasoning, they suggest that ethical thought and discourse must essentially involve the exercise of practical reasoning instead.

In the contemporary metaethical debate, expressivist and constructivist views can be viewed as inspired by these ideas from Hume and Kant respectively.² As I’ll use the labels, *expressivism* is the view that ethical discourse differs semantically from matter-of-factual thought and discourse in that ethical sentences mean what they do not in virtue of representing matters of fact but rather in virtue of conventionally expressing sentiments, attitudes, acceptances of norms, or (if you like) ethical ‘beliefs’ (as long as these are understood as essentially different from matter-of-factual beliefs). By contrast, *constructivism* holds that ethical thought and discourse differ from matter-of-factual thought and discourse in being not about putatively antecedently obtaining states of the world but rather about a special sort of facts which is ‘constructed’ out of the constraints of practical reasoning. Each of these views purports, in its own way, to respect the traditional motivation for irrealism in metaethics stemming from skepticism about ethical properties.

To be sure, contemporary expressivists and contemporary constructivists will each acquiesce in *talk of* ethical properties ‘in a sense’, but they’ll deny that the sense in which they accept such talk puts these properties on a par with matter-of-factual properties, whose instantiation is empirically discoverable through the exercise of perception and theoretical reasoning. In this vein, contemporary expressivists sometimes say that there are ethical properties in a *minimalist* sense of ‘property’. This means, roughly, that there’s a pleonastic equivalence between, e.g., ‘ ϕ ing is wrong’ and ‘ ϕ ing has the property of being wrong’. And, similarly, contemporary constructivists sometimes say that there are ethical

² With expressivism, I have in mind the family of views represented by Ayer (1936/1946), Stevenson (1936), Blackburn (1993, 1998), Gibbard (1990, 2003), Timmons (1999), Ridge (2006, 2007). With constructivism, the view has been less well worked out in the metaethical literature, but it owes inspiration to Rawls (1980), and I mostly have in mind Korsgaard (2003), and Street (2008). There are realist versions of expressivism (Copp 2001) and constructivism (Lynch 2009: ch. 8). These views fall outside the scope of my interest in this paper.

properties in a *constructed* sense of ‘property’, which parallels the constructed sense in which there are ethical facts. This means, roughly, that ethical properties are not part of the fundamental nature of reality but constructed by practical reasoners.

Whether minimalist or constructed properties, both sorts of irrealism agree that ethical properties are not the sort of thing that can be empirically discovered or known to be instantiated through the exercise of theoretical reasoning. And, as such, expressivists and constructivists don’t count them in the final ontological reckoning, even if there is some weaker sense in which they accept them. Moreover, each view picks up on either the Humean or the Kantian understandings of the apparently internal connection between ethics and action.

Expressivism offers easy explanation of the apparently *regular* connection between one’s ethical judgments and one’s motivation to act in their stead. For, if those ethical judgments are thought to play a functional role more like desires than matter-of-factual beliefs, then they can provide what else must be added to one’s matter-of-factual beliefs to generate motivation to act. And, by viewing ethical claims as expressing desire-like states of some sort, expressivism avoids a realist picture according to which these claims are made correct or incorrect by empirically discoverable matters of fact.

Constructivism offers easy explanation of the apparently *rational* connection between ethics and action. For, if ethical properties are not found in the world but constructed out of the constraints of the very sort of reasoning – practical reasoning – which determines which actions are rationally warranted, then appreciation of them can (at least sometimes) provide what more is needed for rationally warranted action, beyond the sorts of matter-of-factual beliefs that are within the providence of theoretical reasoning. And, by viewing ethical properties not as discoverable or always already there in the world but as constructed out of the logical constraints on practical reasoning, the constructivist also claims consonance with naturalistic-driven skepticism about our empirical access to such properties.³

A popular realist response to these contemporary irrealist views is to argue that they are inconsistent with obvious surface-level appearances of ordinary ethical thought and discourse. Casual reflection on ethical thought and discourse reveals very quickly that ordinary folk speak (and so appear to think) as if (i) some ethical claims are true and others are false, (ii) there are ethical facts, (iii) when we disagree about some ethical matter, it’s often because we have differing beliefs about what the ethical facts are, and (iv) in many easy cases, we know quite well what the ethical facts are. Although Hume and Kant, as well as early expressivists and constructivists may not have been too concerned to accommodate these appearances, doing so has come to be an important desideratum for anyone hoping to claim the mantle of irrealism in the contemporary debate.⁴ If one thinks that reality doesn’t include ethical properties in the final ontological reckoning, then one owes an account of (i)-(iv). Otherwise one’s irrealism will commit one to something like a radical error-theory or a pernicious relativism after all.

³ Both expressivists and constructivists will often claim to be able to capture the other sense in which there appears to be an ‘internal’ connection between ethics and action. I don’t deny that this is possible; rather my point here is to sketch the central motivation for the differing views.

⁴ This is mainly because of the pioneering work of Blackburn’s (1984: ch. 6, and 1993) development of the quasi-realist program, which seeks to recapture for irrealists the sorts of semantically infused talk (e.g. of truth, fact, property, etc.) that tempt many to realism in various areas. For more elucidation and further citations, see note 9 below.

My objective in the rest of this paper is to explore some constructivist and expressivist options for responding to this argument. My conclusion will be that, although both constructivists and expressivists can capture the first three features, the fourth (involving the possibility of ethical knowledge) causes special problems for these versions of irrealism. I'll end with some comments about where I think irrealists should begin to look for a response to these special problems, which points somewhat surprisingly towards an alternative inferentialist form of irrealism.

Constructivism

Let's begin with constructivism, because, at first blush, it seems the constructivist should have an easier time responding to the argument from appearances than the expressivist. For the constructivist holds very explicitly that there *are* ethical facts, it's just that these facts are 'constructed' by the logical constraints on practical reasoning rather than found among the antecedently obtaining matters of fact. And, if there are ethical facts, then it's no wonder that ordinary folk speak as if (i) some ethical claims are true and others are false, (ii) there are ethical facts, and (iii) when we disagree about some ethical matter, it's often because we have differing beliefs about what the ethical facts are. But, when we begin to reflect on the fourth surface-level feature of ordinary ethical thought and discourse – viz., (iv) in many easy cases, we know quite well what the ethical facts are – this story begins to look a lot more tenuous than it may at first appear. Or, at least, so I will now argue.

What do constructivists have to say about the nature and possibility of ethical knowledge? Clearly one thing they cannot say is that we have this knowledge just like we have knowledge of ordinary matters of fact. It's part of the whole point of irrealism to deny that the ethical properties of things are discoverable through perception and theoretical reasoning like ordinary matters of fact. But what's the alternative?

Unfortunately, as it turns out, prominent constructivists have said very little. In her paper 'Realism and Constructivism in Twentieth Century Moral Philosophy', which has sparked much of the recent discussion of constructivism in metaethics, Korsgaard does briefly consider the issue of ethical knowledge. But she does so as a way of criticizing her realist competitors who, according to her, make the mistake of thinking of ethics as a pursuit of knowledge, which we can then apply in practice. She writes, 'The moral realist thinks of practical philosophy as an essentially theoretical subject. Its business is to find, or anyway to argue that we can find, some sort of ethical knowledge that we can apply in action' (2003: 118; cp. also 1996: 44). But, besides thinking that this is wrongheaded, apparently Korsgaard thinks that the issue of ethical knowledge is almost entirely irrelevant. She continues: 'According to constructivism, the only piece of knowledge that could be relevant here is knowledge that the problems represented by our normative terms are solvable, and the only way we can find out whether that is so is by trying to solve them' (*ibid.*). It takes a bit of unpacking to understand this curious claim.

Korsgaard (*ibid.*: 104-105) suggests that it is crucial in seeking to understand the nature of ethical thought and discourse that we distinguish between three different questions one might ask about concepts and three correlative dichotomies in theoretical approach to understanding ethical concepts. First, there's a question about the *source* of concepts. Rationalists about a concept say that it is derived from reason, and empiricists say that it is derived from experience. Then, there's a question about the *content* of concepts. Cognitivists about a concept say that its content is constituted by the contribution it makes to propositions, and noncognitivists deny this, holding instead that the content is at least

partially constituted by something else such as the emotive or prescriptive force of its deployment. Finally, she thinks there's a question about the *function* of concepts. Descriptivists about a concept say that its function is to describe things in the world, and constructivists say that its function is to name and think about a solution to a practical problem.

In light of this, Korsgaard's constructivist view of ethical concepts is one about their function rather than about their source or content: 'They are the names of the solutions of problems, problems to which we give names to mark them out as objects for practical thought. The role of the concept of the right, say, is to guide action; the role of the concept of the good might be to guide our choice among options, or of ends' (2003: 116). And it is *conceptions* of the right or the good that are proposed as possible solutions. That is to say that these ethical concepts function to name whatever conception(s) solve the problems for which the concepts are needed. Thus, with Kant and Rawls in mind, Korsgaard writes, 'The task of practical philosophy is to move from concepts to conceptions, by constructing an account of the problem reflected in the concept that will point the way to a conception that solves the problem' (ibid.).

I think it's because she thinks that ethical concepts are the names of solutions to problems that she thinks that the only knowledge that could be relevant to the concerns of ethics is the knowledge that the problems we face are solvable. This is because the point of ethics as a discipline is not to acquire theoretical knowledge of some antecedently existing matter of fact but to generate solutions to practical problems in the form of conceptions of the good, the right, etc., which if successful will constitute constructions of the relevant ethical facts.

Even if that is right, however, the metaethical realist may reasonably claim that Korsgaard has illicitly changed the subject.⁵ His concern is not primarily ethical but *metaethical*; and he demands an explanation of how constructivism squares with the ordinary practice of claiming and ascribing ethical knowledge. When someone says 'I know that what I did was wrong' or 'We all know that one should help those most in need,' surely it's implausible to think that this knowledge is simply knowledge that a solution to the practical problem of what to do or not to do in certain circumstances is possible. Yes, we wouldn't claim and ascribe such knowledge unless we thought (at least implicitly) that the practical problem was solvable, but these knowledge attributions seem to be more. They seem to be claims to know not only *that there is a solution* but also *what the solution is*.

Perhaps, however, this line of objection is not maximally charitable to Korsgaard. After all, she does hold that there are ethical facts, albeit constructed facts. So, maybe a more plausible response to the realist's objection would come from a fuller idea of what's involved in a fact's being constructed and how we can then know such facts. I'm not entirely sure what metaethical constructivists mean by 'constructed', but Korsgaard gives us an example having to do with human artifacts such as chairs. She writes, 'Why do we have the concept of chair? Certainly not because it would form part of the absolute conception of the world...We have the concept of "chair" because the physical construction of human beings makes it possible, and occasionally necessary, to sit down' (ibid.: 117). I take it the idea is that facts about chairs are constructed when humans think of and make, i.e. *construct* chairs, as a solution to our need to sit down. And it certainly seems like we can know many things about chairs. For example, we know that chairs generally have four legs. Thus, if ethical facts are constructed like facts about chairs are

⁵ Compare FitzPatrick (2005) and Hussain and Shah (2006) where a similar point is made.

constructed, the ethical constructivist should have a complete answer to the realist argument from appearances.

However, this line of thought is also problematic. If we're thinking that facts about chairs are 'constructed' – and thus shouldn't figure in the final ontological reckoning – because the concepts with which they are cognized would not figure in the 'absolute conception of the world', then surely that doesn't mean that they are any less real than facts about things like rocks, which we are 'found in nature'. This is because, Korsgaard may be right that the concept of chair won't form part of the 'absolute conception of the world', but neither will the concept of rock. This is simply because both kinds of things can be conceived of in different and more objective ways. For example, if the conceptual scheme of physics approximates the absolute conception of the world better than the conceptual scheme of commonsense, which includes the concepts of 'chair' and 'rock', then what commonsense thinks of as chairs and rocks may also be thought of by physics as collections of physical particles arranged in particular ways.

If, however, we're thinking that facts about chairs are 'constructed' – and thus shouldn't figure in the final ontological reckoning – because their existence *qua* chair, depends on a conceptual framework involving our practical need to sit down, this may explain their ontological difference from facts about rocks, but it doesn't make for a useful model for ethical knowledge. For most of what we know about chairs (e.g. that they generally have four legs) we know through perception and theoretical reasoning (including induction). However, the point of the original irrealist argument stemming from Hume and Kant was precisely that we don't seem to acquire ethical knowledge in this way. This means that the metaethical constructivist cannot appeal to the fact that we know a lot about constructed things like chairs to explain how we know about ethical facts, which she also conceives of as constructed. Indeed, it's not even clear that we should think of these two kinds of fact as constructed in the same sense of 'constructed'.

So, before the constructivist can fully meet the realist argument from appearances, she needs a better account of the sense in which ethical facts are ostensibly constructed, which can then be used to develop a better account of ordinary attributions of ethical knowledge. Street (2008) points us in what may seem like a more fruitful direction with her discussion of the possibility of error in normative judgment. She writes, 'While there are, ultimately, no normative truths that hold independently of our evaluative attitudes — while normative realism is false — it does not follow that it's impossible to go wrong with one's normative judgments' (*ibid.*: 207). The reason she thinks this doesn't follow is that her constructivist view has it that 'the fact that X is a reason to Y for agent A is constituted by the fact that the judgment that X is a reason to Y (for A) withstands scrutiny from the standpoint of A's other judgments about reasons' (*ibid.*: 223). And, given the (constructed) existence of such facts, it seems possible to know them to obtain.

What would that involve? Street doesn't say very much about the possibility of ethical knowledge, but she does have a worked out view of what it is for an ethical judgment to be true. According to her, truth in ethical judgment is constituted by the judgment's ability to withstand scrutiny from the standpoint of other normative judgments. More specifically, because she wants to avoid the relativist conclusion that normative judgment is radically relative to each speaker's point of view, she claims that the truth of such judgment depends on its ability to withstand scrutiny from the standpoint of the implicated agent. Moreover, the idea is not, I take it, that truth in ethical judgment comes only when the judgment stands up to *actual* scrutiny from the standpoint of other judgments of the implicated agent; rather, the idea is that truth comes when a judgment *would* withstand such scrutiny. And, if

that's the right interpretation, Street's constructivism comes to the view that ethical facts are constructed out of the ability that ethical judgments of them have to withstand scrutiny from the point of view of the implicated agent. This makes ethical facts ontologically dependent on human judgments/attitudes in a way that facts about chairs and rocks generally are not. Our judgments that chairs generally have four legs or that rocks are hard are not true in virtue of their ability to withstand scrutiny from other judgments – at least that's a standard motif of realism. Street thinks the opposite is the case for ethical judgments, which means that she avoids one pitfall of Korsgaard's view.

If that's right, we should be able to generate an account of ethical knowledge out of Street's constructivist account of true ethical judgment. What's needed is an account of what makes true ethical judgment (as Street is conceiving of them) into knowledge. Epistemologists call this *warrant*. So what Street needs to make sense of ethical knowledge is an account of epistemic warrant for ethical judgments.

To this end, one thing she might try is to simply deploy one of the accounts of warrant already widely discussed in epistemology. Perhaps – to pick a representative instance of epistemic externalism and epistemic internalism – ethical knowledge requires true judgment formed by a reliable process or true judgment that is appropriately based on evidence.

There is, however, a problem with this for Street. For, recall that she thinks the truth of ethical judgments consists in their ability to withstand scrutiny from the standpoint of the normative judgments of the implicated agent. That is, it's an important feature of Street's view that the relevant standpoint of scrutiny for a judgment like

(1) Amy's judgment that Bush ought not to have taken the country to war

is the standpoint comprising the normative judgments *of the implied agent*, which in this case is Bush. This is how she avoids a radical form of speaker-relativism about the truth-conditions for judgments like the one described in (1). However, if that's right, then for (1) to be warranted, it would seem that certain things would have to be true about the relationship between Amy and the standpoint comprising Bush's normative judgments. For instance, on a reliabilist conception of warrant, in order to be knowledge, (1) would have to be formed by a process that reliably forms judgments that withstand scrutiny from Bush's standpoint. But it's implausible to think that all those who know that Bush ought not to have taken the country to war can reliably form judgments that withstand scrutiny from Bush's perspective; Bush's system of normative judgments may be entirely inscrutable to them. But, on an evidentialist conception of warrant, by contrast, in order to be knowledge, (1) would have to be appropriately based on evidence that it would withstand scrutiny from Bush's standpoint. However, again, it's implausible that all those who know that Bush ought not to have taken the country to war have based their judgment on evidence that it would withstand scrutiny from Bush's standpoint, which may (to repeat) be rather inscrutable.

The example of (1) and the reliabilist and evidentialist conceptions of warrant is just an example. However, I think it represents a real problem for Street's brand of constructivism to capture ordinary talk of ethical knowledge. For even if she denies that example is a plausible example of ethical knowledge, it's surely part of ordinary ethical thought and discourse to make ethical judgments which implicate agents whose system of normative judgments is inscrutable to us. Indeed, often we make generalized judgments about, e.g., what is right/wrong for *everyone* or what *anyone* ought to do in some situation, and it is implausible to think that we have judgment-forming processes that reliably track or that we base these judgments on evidence about what would withstand scrutiny from *everyone's* or *anyone's* system of normative judgments.

Now, admittedly, the reliabilist and evidentialist conceptions of epistemic warrant used in raising the objection just mooted are only two instances of possible accounts of epistemic warrant. However, it's hard to see what alternative conception will avoid the problem I've raised. The only alternative that suggests itself (to me, at least) is to claim that ethical judgments are special in that, for them, warrant and truth come to the same thing. For Street's brand of constructivism, that would mean that a judgment that someone ought to do something is both true and warranted in virtue of its ability to withstand scrutiny from the standpoint of the implicated agent's normative judgments. This would save constructivism's ability to accommodate ordinary talk of ethical knowledge. However, it seems to me to flounder on the problem of wrongly counting certain kinds of luckily true judgments as knowledge.

Luckily true judgments don't seem to be knowledge.⁶ The post-Gettier literature has taught us that even judgments that are warranted on a reliabilist or evidentialist conception of epistemic warrant could be true primarily because of luck and so not knowledge. But this would also be the case with a class of supposedly special judgments for which both truth and warrant amount to withstanding scrutiny from the standpoint of some other class of judgments. To appreciate this, consider again the example of (1). Now, however, contrast that with

(2) Joe's judgment that Bush ought to have taken the country to war, where both Amy and Joe don't know very much about Bush's system of normative judgments. Perhaps they have just guessed, or perhaps they have made their judgments by projecting their own general ethical views onto Bush. Now, assuming that there's a fact of the matter about which of these judgments would withstand scrutiny from the perspective of Bush's normative judgments, either Amy or Joe will have judged truly and the other falsely, according to Street's constructivism. But, given her view of what constitutes the truth of ethical judgments, it seems that, whichever person has judged truly, he or she has done so primarily because of luck. So, unless we want to count ethical judgments that are true primarily because of luck as knowledge, it will do no good in attempting to accommodate the possibility of ethical knowledge for Street's constructivist to say that truth and warrant come to the same thing in the special case of ethical knowledge.

Perhaps Street could respond by arguing that we don't have very much ethical knowledge after all. She doesn't need to eliminate all ethical knowledge. In cases, e.g., where the implicated agent is the person making the judgment, it may be easier to have warranted beliefs about what would or would not withstand scrutiny from the standpoint of the agent's other normative judgments. However, in third-personal cases, Street could claim that it will be more difficult to have warranted ethical judgments, and ethical knowledge will simply be much harder to come by.

I think this would represent a coherent irrealist stance on the possibility of ethical knowledge. But it wouldn't answer the realist's argument from surface-level appearances that I am considering in this paper. For recall that the realist's argument emerges in a debate between realists and irrealists who want to avoid some radical error-theory or pernicious relativism in our metaethics. Originally, constructivism seemed to offer some promise of such irrealism, since constructivists recognize that there are, in a sense, ethical

⁶ This claim requires some qualification. For some kinds of luck are compatible with knowledge. For example, knowledge that a pelican just flew by the car may depend on being lucky enough to have looked out the window at the opportune moment. However, other kinds of luck are not compatible. I'm discussing those other kinds of luck here (without attempting to specify them precisely). See Pritchard (2005: ch. 5-6) for discussion of permissible and impermissible forms of epistemic luck.

truths, facts, and disagreements in belief. This helps them to accommodate ordinary talk of such things. However, what I think we should conclude from the preceding discussion is that metaethical constructivists like Korsgaard and Street have not yet⁷ offered us a sense in which there can be ethical knowledge on their view; and until they do so, they won't have answered the realist's argument from appearances.

Expressivism

Recall that the argument we are considering against irrealist views like constructivism and expressivism is that they cannot capture the surface-level appearances of ordinary ethical thought and discourse; this includes ordinary folk's propensity to speak as if there are ethical truths, facts, disagreements in belief, and knowledge. It may seem that expressivists have no chance of deflecting this argument, since, inspired by Hume, they think that ethical thought and discourse does not mainly involve beliefs in matters of fact but rather coordinating our sentiments in order to live well in community. Indeed that view is supposed to found an alternative semantics for ethical discourse, according to which ethical sentences get their meaning not from what they (purport) to represent but from the desire-like mental states they (conventionally) express. Recently, however, expressivists have been resourceful in accommodating the appearances of ordinary ethical thought and discourse. They typically do so by interpreting the surface-level features of ordinary ethical thought and discourse in a way that is neutral between realism and irrealism, thus insisting that the appearances do not sway the debate one way or another.

For example, when it comes to truth, expressivists now typically embrace some general account of truth that is neutral on the ontology of morality – such as deflationism or minimalism.⁸ Then they claim that adopting such a theory of truth allows us to say that there is no more ontological commitment involved in saying, e.g., that it's true that torture is wrong than in saying that torture is wrong. To be sure, that doesn't entail that there's *no* ontological commitment involved. But what it does mean is that the mere fact that ordinary folk think and talk as if moral claims are true is no threat to the expressivist. We might call this an 'accommodation story' for ordinary talk of ethical truths.⁹ With it, the expressivist can accommodate this feature of ordinary ethical thought and discourse.

The accommodation story for ordinary talk of ethical facts and ethical beliefs is very similar and intimately related. In ordinary ethical thought and discourse, it seems that the notion of 'fact' can be replaced without much loss by talking about truth. 'It's a fact that torture is wrong' seems not to be appreciably different in ordinary discourse from 'It's true that torture is wrong'. So, as long as the accommodation story for truth is cogent, then it will extend easily to facts. Likewise, 'She believes that torture is wrong' isn't appreciably different in ordinary discourse from 'She thinks it's true that torture is wrong'. So, again, if the accommodation story for truth works, then there shouldn't be a problem for the expressivist to recognize a category of ethical beliefs, even though he denies that these

⁷ There's a promising footnote in Street's paper where she hints at a constructivist metaepistemology, but she doesn't develop the view there. It's something like this which I think a constructivist will need to overcome the realist argument from appearances.

⁸ For example, one of the views defended by Quine (1970), Field (1986, 1994), Horwich (1990), or Wright (1992).

⁹ This strategy received its first explicit expression when Blackburn proposed the 'enterprise of quasi-realism,' which 'tries to earn, on the slender basis [of expressivist anti-realism], the features of moral language...which might tempt people to realism' (1984: 171). It has been subsequently developed by *inter alia* Blackburn (1992), Stoljar (1993), Horgan and Timmons (1993, 2000), Timmons (1999), and Gibbard (2003).

beliefs are matter-of-factual beliefs, since they have a different psycho-functional role and direction of fit with the world.¹⁰ And the expressivist will construe ostensible disagreement in ethical beliefs as a different sort of disagreement from disagreement over ordinary matters of fact. Philosophers may want to make fine distinctions between how they use notions like truth and fact, or belief and thinking true, but ordinary thought and discourse seems to run these notions together fluidly. And this means that no mere appeal to ordinary talk of ethical truth, fact, or belief will undermine expressivism – at least not if the expressivist is allowed his favored interpretation of these terms.

This last ‘if’ is a big one, of course, but assuming we accept it, I think it puts the expressivist very quickly on a par with the constructivist with respect to the realist argument from appearances. If we reject representationalist assumptions behind realism, it turns out to be relatively easy to tell a constructivist or expressivist story that makes sense of the way ordinary ethical thought and discourse deploys the notions of truth, fact, and belief. As with constructivism before, however, I think the notion of knowledge poses a special problem for the expressivist. This is because knowledge seems to require not only true judgment but also some sort of warrant for this judgment. And, as we saw above, warrant seems like it must involve some element that avoids the possibility of luckily true judgment. However, as irrealists, expressivists cannot make the appeal epistemologists typically make to the idea of reliably tracking the properties of things in the world. For, although they can talk with the ordinary folk and speak of the fact that, e.g., torture is wrong, they don’t want to understand this talk in a way which involves treating our knowledge of it as empirically tracking some property of wrongness, which figures in the final ontological reckoning. That means that the expressivist owes an alternative account of the role of the notion of knowledge in ordinary ethical thought and discourse.

Fortunately, expressivists have said significantly more about the ordinary talk of ethical knowledge than their constructivist compatriots. In pursuit of an accommodation story here, they seek to embrace some general account of knowledge that is neutral on the metaethical debate. For instance, Blackburn suggests that, ‘...the primary function of talking of “knowledge” is to indicate that a judgment is beyond revision. That is, we rule out any chance that an improvement might occur, that would properly lead to a revision of the judgment’(1998: 318). And Gibbard claims that knowledge attributions, such as ‘Joe knows there are cows on the hill...means very roughly...that judgments like his are to be relied on.’(2003: 227).

This maneuver handles the problem with luckily true ethical judgments, which seemed to trip up constructivists like Street. On Blackburn’s version of the idea, he can say that our intuition that such judgments are not knowledge is just a reflection of our view that an improvement might occur, i.e. that the luckiness might be revealed, which would properly lead to revision. And on Gibbard’s version of the idea, he can say that these intuitions are a reflection of our view judgments that are formed in a way that is subject to luck are not one’s that are to be relied on.

However, this maneuver does raise an important question: What is going on when we indicate that a judgment is ‘beyond revision’ or ‘to be relied on’? Are we stating a matter of fact or are we expressing some desire-like state? I take it both Blackburn and Gibbard mean the latter. That is, Blackburn’s idea is that, when we move from the fact that

¹⁰ Perhaps ordinary speakers are committed to the in principle resolvability of disagreements in ethical belief.(Smith 1994: 5-6) However, I doubt this is true. It’s a different issue whether expressivism is consistent with a defensible *theory* of ethical disagreement. On this, Egan (2007) poses a serious challenge to expressivism which is outside the scope of the present paper.

someone judges that *p* to claiming that he knows that *p*, we are expressing an epistemic stance towards this person's judgment. We are marking it out as 'beyond revision' in the sense that 'no further useful investigation or thought ought to undermine the [judgment]' (1996: 87). Gibbard is even clearer about embracing epistemic expressivism. His technical terms for judgments that require an expressivistic analysis is 'plans' and 'plan-laden', and he writes, '...attributions of knowledge are plan-laden...Concluding that Joe knows, then, amounts to planning to rely on his judgment' (Ibid.).

However, we might worry that someone's judgment could be thought to 'beyond revision' in Blackburn's sense or 'to be relied on' in Gibbard's sense but nonetheless false; and surely false judgments do not count as knowledge. Perhaps Blackburn and Gibbard could avoid this problem with a fuller account of what sorts of judgments are 'beyond revision' or 'to be relied on', but a simple way to avoid the problem is to supplement Blackburn's and Gibbard's accounts of 'knows' with a truth condition. The idea would be that when we ascribe knowledge, we are not only expressing an epistemic stance towards someone's judgment that *p*, we're also endorsing *p* as true.¹¹

In sum, the expressivist response to the problem of accommodating ordinary talk of ethical knowledge is to extend the scope of the expressivist theory to include epistemic judgments like knowledge attributions alongside ethical judgments. The most plausible version of expressivism for knowledge attributions is a form of what is sometimes called 'hybrid' expressivism, because it holds that knowledge attributions have some matter-of-factual content (i.e. that the subject has a judgment which we regard as true) but they also have some extra content that is explicable in only in terms of the expression of a desire-like state that might be characterized as 'indicating that the judgment is beyond revision' or 'planning to rely on the judgment'.

The attraction of this position for the expressivist about ethical judgments is parallel to the attraction of a minimalist or deflationist account of truth. It allows the expressivist to interpret ordinary talk of ethical knowledge in a way that is neutral on the metaethical debate. If we don't, in general, understand knowledge attributions in terms of the putative knower reliably tracking some empirically discoverable matter-of-factual properties, then there is theoretical space to hold that ethical judgments can be coherently called knowledge, even though they don't involve such tracking of the external world. We'll call them knowledge when we think they are true and meet the epistemic norms we accept. As Gibbard puts the point, '...knowledge of what to do...goes beyond just happening to track the right property in our planning...we fasten onto the right things to live for because we are in a condition to be trusted on such matters – or so we judge' (2003: 235).

This accommodation story overcomes some of the problems for constructivism stemming from the realist's argument from appearances, and it does so in a way that seems to me to be consistent with ethical irrealism and the advantages for explaining the internal connection between ethical matters and action (at least on the Humean interpretation of this connection). However, extending expressivism to cover knowledge attributions as well as ethical claims highlights two problems that have led me, anyway, to look for other options.

The first problem doesn't really have anything to do with *epistemic* expressivism per se, but I think extending the expressivist's semantic view from ethical sentences to epistemic sentences may highlight for some the way in which expressivism involves a highly controversial semantic view. Unlike constructivists and realists, who can endorse a standard

¹¹ This is the basic strategy I pursued in my own previous defense of epistemic expressivism; see my (2007).

account of the meaning of ethical sentences in terms of their truth-conditions (even while they disagree about what constitutes these truth-conditions), expressivists found their irrationalism on an alternative semantic account of ethical sentences, which now looks like it has to be extended to epistemic sentence, in order to answer the realist argument from appearances. The expressivist holds, roughly, that these sentences mean what they do not in virtue of their truth-conditions but in virtue of the desire-like state that they express. Given the apparent discontinuities in function and practical import between ethical language and paradigmatically descriptive language, attempting to work this view into a viable semantics may seem like a worthwhile project. However, it involves theoretical costs associated both with attempting to give a non-circular 'ideationalist' or 'psychologistic' account of semantic content¹² and with capturing the logical continuities between expressivistic language and matter-of-factual language.¹³ Thus, when the expressivist is required to extend that project to epistemic sentences as well, one may begin to wonder whether the theoretical costs it incurs don't outweigh whatever advantages it purports to bring. Another way to put the point is that if the expressivist semantics for ethical sentences requires an expressivist semantics for epistemic sentences, then Frege-Geach type worries are even more pressing than they already seemed in the ethical case.

Expressivists can legitimately ignore this first problem, since they are already committed to there being an answer to any objections generated by their alternative semantics for ethical and epistemic sentences. However, the second problem I see challenges the overall cogency of the expressivist position, once epistemic expressivism is added to ethical expressivism in an attempt to accommodate appearances.

To appreciate this problem, recall that part of the original motivation for irrationalism in metaethics is the apparently special link between ethics and action, though, as we saw above, Humeans and Kantians have different accounts about the nature of this link. Roughly, Humeans think that it is a matter of a *regular* connection between ethical judgments and being suitably motivated. In contrast, Kantians think that the link is a matter of a *rational* link between ethical judgments and acting in certain ways. Expressivists follow Hume on this, which is why they claim that ethical judgments mainly function to help us to coordinate sentiments in pursuit of communal cooperation and fellow feeling. Thus, a large part of the motivation for expressivism in metaethics is a commitment to the Humean account of the special internal connection between ethics and the will combined with the claim that expressivism best explains/captures this connection.

However, I think it generates a dilemma if expressivists now want to extend their expressivism to cover epistemic judgments as well as ethical judgments. The dilemma is that either a parallel argument from the Humean interpretation of internalism is also a powerful motivation for expressivism about epistemic judgments or this sort of argument is not a very powerful motivation for expressivism about ethical judgments in the first place.

The problem along the first horn is that we don't have the same sorts of internalist intuitions when it comes to epistemic judgments such as knowledge attributions as we do

¹² The sort suggested by Grice (1957), and worked out in different ways by Schiffer (1972) and Davis (2003).

¹³ No participant to the metaethical debate denies that it is highly debatable whether expressivism about ethical claims can provide a semantics for ethical sentences that is even structurally adequate for the task of distinguishing the semantic values of ethical claims under simple logical operations and propositional attitude reports. Recently Schroeder (2008), picking up some strands from *inter alia* Unwin (1999) and Dreier (2006), has persuasively argued that most extant versions of expressivism have inadequate explanations of the semantic function of negation and propositional attitude reports.

with ethical judgments. With ethical judgments, the internalist intuition (interpreted Hume's way) was that, if someone judges, e.g., that she ought to give to charity, then *ceteris paribus* it'd be strange if she weren't at least somewhat motivated to give to charity. As we said, the motivation seems to attach *regularly* to the ethical judgment. That's pretty much accepted as an intuitive datum in the metaethical attempt to explain ethical judgments. But there doesn't seem to be a parallel for epistemic judgments. If someone judges, for example, that he knows that the moon is about a quarter of a million miles from the earth, then what do we expect him to be motivated to do? It's not clear that there is anything in particular.

An expressivist might suggest that this person will be motivated to believe that the moon is about a quarter of a million miles away from the earth. But that's an odd thing to say for two reasons: First, unlike the previous example of giving to charity, which is yet to be done, when one judges that he knows that the moon is about a quarter of a million miles from the earth, one typically already believes this to be true. Second, believing isn't (usually) the sort of thing that can be the object of our motivations, since believing is (usually) involuntary.¹⁴ Alternatively, an expressivist might suggest that someone who judges himself to have such knowledge will be motivated to testify that the moon is about a quarter of a million miles from the earth, if asked. But, even if that is indeed true, it's not clear that the epistemic judgment is what bears the internal connection to this motivation; after all, one who firmly believes that the moon is about a quarter of a million miles away (but has never considered whether he knows this to be true) would be similarly motivated to testify.

Something else an expressivist might try is to suggest that someone who judges himself to know that *p* will be motivated to act as if *p* is true. However, that's such a weak requirement, that I think it will prove almost impossible to form any plausible intuition about whether there is a regular connection between self-attributions of knowledge and such motivation. For example, if I wanted to test whether our knower is indeed motivated to act as if the moon is about a quarter of a million miles from the earth, how would I go about doing so? What is it to act as if the moon is about a quarter of a million miles from the earth? How would I distinguish this ostensible motivation from other very similar possible motivations, such as acting as if the moon is about 380,000 kilometers from the earth? I don't think these questions are easy for the epistemic expressivist to answer in a plausible way. The trouble here is not that there aren't some motivations that might attach somewhat regularly to self-attributions of knowledge, but that there are enough other things that could explain these motivations that the internalist intuition in the epistemic case must be much weaker than in the ethical case.

Even if there is no regular connection between first-personal self-attributions of knowledge and being motivated to act in certain ways that could serve in a parallel argument from internalism to expressivism, the expressivist could try to find the regular link between knowledge attributions and the will in third-personal attributions of knowledge to others. For example, Gibbard's suggestion seems to be that claiming that someone knows that *p* amounts to planning to rely on him with respect to *p*. So, maybe the internal connection is to being motivated to rely on someone's judgment. However, that can't be right given the truth requirement on knowledge. Because of this, one wouldn't ascribe knowledge that *p* to someone *S* unless one already took *p* to be true. This suggests that the elusive internal connection might be found between claims of the form '*S* knows

¹⁴ For more on involuntarism and why it's nevertheless compatible with normative epistemic claims, see my (2008a).

whether p and being motivated to rely on with respect to p . However, that too fails to be plausible. I may be able to tell you which of my students know whether Ayer was an expressivist or a subjectivist, but that doesn't mean that I'd be motivated to rely on their judgments on this matter – to put this somewhat arrogantly: I don't *need* to rely on their judgments.

There may yet be some internal connection between some epistemic judgment and motivation to action, but, even if there is, its presence surely isn't an intuitive datum of the metaepistemological debate like the presence of the connection in the ethical case is. Thus, if my dilemma is correct, then the expressivist will be forced onto the second horn, which involved conceding the weakness of the argument from internalism to expressivism. Although being forced onto this horn of the dilemma surely wouldn't be the end of expressivism – there may, after all, be other arguments for the view beside the argument from internalism – I think it clearly represents a significant theoretical cost for the expressivist.

So, although I think expressivists are on a par with constructivists with respect to accommodating the ordinary use of notions of ethical truth, fact, and belief, the accommodation story they give about ordinary talk of ethical knowledge raises separate problems for the theory. And these problems will remain fodder for the realist's argument from surface-level appearances.

Conclusion

I've been exploring constructivist and expressivist responses to the realist argument from the surface-level appearances of ordinary ethical thought and discourse. I've argued that the feature of those appearances that causes the most trouble for constructivists and expressivists is the way ordinary folk talk of ethical knowledge. In a way, this shouldn't be surprising, since part of the original motivation for irrealism was skepticism about empirical access to the ethical properties of things. Does this mean that irrealism is destined to be undermined by the realist argument from appearances? I'm not sure, but I think the discussion of constructivism and expressivism may point the way to an alternative response to the argument than the ones we've already seen. I'll conclude with some speculative comments about how that might go.

The expressivist's basic maneuver for capturing ordinary talk about ethical knowledge is to extend his form of irrealism to apply to epistemic thought and discourse just like it applies to ethical thought and discourse. However, one of the main motivations for expressivism in the ethical case – the argument from Humean internalism – seems like it doesn't carry over to the epistemic case. This problem threatens epistemic expressivism, but it doesn't necessarily threaten epistemic irrealism. For, of course, there are other irrealist options. One of these would be to develop a constructivist form of epistemic irrealism. Perhaps this could pick up on the Kantian interpretation of internalism rather than the Humean interpretation. Street hints that this is the direction she'd be inclined to go, writing in a footnote: 'While I believe that constructivism also provides the best account of the normativity of epistemic reasons, I will not discuss that here' (2008: 209). It will be key in working out this idea to explain the way epistemic reasons for matter-of-factual beliefs differ from epistemic reasons for ethical beliefs, while retaining the idea that neither sort of belief can count as knowledge if it is merely luckily true.

However, I think another option lies in the inferentialist idea articulated by Sellars in his slogan that, '...in characterizing an episode or state as that of knowing, we are not giving an empirical description of that episode or state; we are placing it in the logical space

of reasons, of justifying and being able to justify what one says'(1956). Fleshing out the full implications of this idea is beyond the scope of the present paper, but three parting comments are in order.¹⁵

First, I take Sellars idea of 'placing an episode or state in the logical space of reasons' to be both interestingly similar to and importantly different from Blackburn's idea of 'indicating that a judgment is beyond revision' and Gibbard's idea of 'planning to rely on someone's judgment'. These ideas align in the naturalist-inspired focus on the subject of epistemic thought and discourse rather than its putative object. That is Sellars would agree with Blackburn and Gibbard that epistemic notions allow us to take up a distinctive epistemic stance. This is, I take it, why Sellars denies that characterizations of an episode or state as one of knowing are 'empirical descriptions'.

However, the sort of view of epistemic claims which results from Blackburn and Gibbard's ideas would, insofar as it parallels their view of ethical claims, be one where the relevant epistemic stance is psychologically articulated. That is, it's explained in terms of its causal-functional role in the psychology of the person who manifests it. It's here that expressivists seek to draw an interesting disanalogy between the target area of thought and discourse (epistemic or ethical) and matter-of-factual thought and discourse. This is why they speak of this state as ultimately more desire-like than belief-like. And expressivism holds that the meaning of epistemic claims is determined in part by their role as expressions of that state.

In contrast, someone following out Sellars idea would seek to articulate the state in terms of its inferential role, which in turn would be explained in terms of the inferential potential of the relevant concepts in the community of concept-users who have epistemic concepts. It is here that he would seek to draw an interesting disanalogy between epistemic thought and discourse and matter-of-factual thought and discourse. The idea, roughly, would be to argue that the former is essentially caught up in governing the practice of giving and asking for reasons, while the latter is essentially caught up only in theoretical inferences towards doxastic states like belief. And this view is naturally allied with an inferentialist semantics rather than an expressivist semantics, since it will hold that the meaning of epistemic claims is determined by their socially articulated inferential role rather than their psychologically articulated expressive role.¹⁶ Or, another way to put this contrast is in terms of the different things expressivists and inferentialists think epistemic terms make it possible for us to do. Expressivists think that the ability made possible by epistemic claims is at its core an ability to express a certain kind of attitude. Inferentialists, by contrast, think the ability made possible by epistemic claims is at its core an ability to have that very attitude with its attendant inferential role.¹⁷

My second parting comment is that I think the inferentialist view of epistemic judgments suggested by Sellars provides a nice way to accommodate ordinary folk's talk of ethical knowledge within an irrealist framework. The key will be the claim that ethical judgments can be 'placed in the space of reasons' just as much as any other sort of judgment. And here both irrealists and realists alike should agree that ethical judgments are ones which we can sometimes justify, even if they disagree about whether that justification

¹⁵ See my (ms) for an attempt at an account of propositional knowledge which is inspired by Sellars' slogan and aims to meet a number of desiderata stemming from current debates in epistemology. See my (forthcoming) for some more discussion of the relationship between epistemic expressivism and epistemic inferentialism.

¹⁶ The sort of semantic program developed by Rosenberg (1974) and Brandom (1994).

¹⁷ Compare Craig (1990), Williams (1992), and Rosenberg (2002: ch. 5-6)

involves empirically tracking some external moral properties which figure in the final ontological reckoning. So, Sellars' inferentialism about knowledge attributions seems to provide a neutral way to accommodate the final feature of ordinary ethical discourse that figures in the realist's argument from surface-level appearances. This means that, if I'm right that the notion of ethical knowledge poses a special problem for irrealists like constructivists and expressivists, then Sellars' inferentialist and irrealist view of knowledge attributions should seem like an attractive way to solve the problem.¹⁸

Third, now that we've seen the beginnings of an irrealist account of epistemic judgments that is neither a form of epistemic constructivism nor a form of epistemic expressivism, but rather a form of inferentialism, I think we should ask whether a similar account of ethical judgments might not provide an attractive alternative form of *ethical* irrealism. That is, might it not make sense to say that in characterizing an action as one that someone ethically should or shouldn't do, we are not (to modify the metaphor from Sellars) giving an empirical description of that action, rather we are placing it in the space of reasons, of justifying and being able to justify what one *does*? If we can inferentially articulate an interesting disanalogy between ethical concepts and the empirical concepts involved in matter-of-factual judgments, we'd have a third irrealist alternative to place beside constructivism and expressivism in the menu of general metaethical theories. The core similarity with these views would be the idea that what we're doing in making ethical judgments is not to describe empirically discoverable matters of fact. The core difference from these views would be the specific view about what ethical terms make possible for us to do; the idea is not that they make it possible for us to *construct* practical facts or to *express* desire-like attitudes with a distinctive pressure on motivation but to *articulate* the difference between good and bad practical inferences.¹⁹

Moreover, although I haven't argued for this claim here, it seems to me that this alternative would better meet the realist argument from appearances. This is mostly because it can accommodate ordinary talk of ethical truth, fact, and belief along similar lines to the expressivist or constructivist while offering a better account of ordinary talk of ethical knowledge.²⁰

Works Cited

- Ayer, A. J. (1936/1946) *Language, Truth and Logic*, London: V. Gollancz Ltd.
- Blackburn, Simon (1984) *Spreading the Word: Groundings in the Philosophy of Language*, Oxford; New York: Clarendon Press; 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 (ed.), *Moral Knowledge: New Readings in Moral Epistemology*, Oxford; New York: Oxford University Press.
- Blackburn, Simon (1998) *Ruling Passions : A Theory of Practical Reasoning*, Oxford; New York:

¹⁸ This does raise the question of how to distinguish realism from irrealism in a given domain, if not in terms of truth, fact, belief, or knowledge. I address this question in my (2008b).

¹⁹ Compare Brandom (2001: ch. 2) for the beginning of such a view. Things get more complicated when we consider ought-claims attaching, in the first instance, not to actions but to states of things.

²⁰ For helpful comments on previous versions of this material, I'm grateful to Michael Ridge, Ana Barandalla Ajona, the Epistemology @ Edinburgh research group, and the participants of the Constructivism and Normative Epistemology workshop organized by James Lenman at the University of Sheffield.

- Clarendon Press; Oxford University Press.
- 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.
- Chrisman, Matthew (2007) 'From Epistemic Contextualism to Epistemic Expressivism', *Philosophical Studies* 135: 225-254.
- Chrisman, Matthew (2008a) 'Ought to Believe', *Journal of Philosophy* CV: 346-370.
- Chrisman, Matthew (2008b) 'Expressivism, Inferentialism, and Saving the Debate', *Philosophy and Phenomenological Research* 77: 334-358.
- Chrisman, Matthew (forthcoming) 'From Epistemic Expressivism to Epistemic Inferentialism', in Haddock, Pritchard, and Millar, eds. *Social Epistemology*, Oxford; New York; Oxford University Press.
- Chrisman, Matthew (ms) 'Knowledge as True Belief One Ought to Have'.
- 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.
- Davis, Wayne (2003) *Meaning, Expression and Thought*, Cambridge: Cambridge University Press.
- Dreier, James (2006) 'Negation for Expressivists', in Shafer-Landau, Russ (ed.), *Oxford Studies in Metaethics, Vol. 1*, Oxford; New York: Clarendon Press.
- Egan, Andy (2006) 'Quasi-Realism and Fundamental Moral Error', *Australasian Journal of Philosophy* 85: 205-219.
- Field, Hartry (1986) 'The Deflationary Conception of 'Truth'', in MacDonald, G. and Wright, C. (ed.), *Fact, Science and Morality*, Oxford: Blackwell.
- Field, Hartry (1994) 'Disquotational Truth and Factually Defective Discourse', *Philosophical Review*, 103 (3): 405-52.
- FitzPatrick, William (2005) 'The Practical Turn in Ethical Theory: Korsgaard's Constructivism, Realism, and the Nature of Normativity', *Ethics* 115: 651-691.
- 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.
- Horwich, Paul (1990) *Truth*, Oxford ; Cambridge, Mass.: B. Blackwell.
- Horgan, Terrence and Mark Timmons (1993) 'Metaphysical Naturalism, Semantic Normativity, and Meta-Semantic Irrealism', *Philosophical Issues* 4: 180-203.
- Horgan, Terrence and Mark Timmons (2000) 'Nondescriptivist Cognitivism: Framework for a New Metaethic ', *Philosophical Papers* 29: 121-153.
- Hussain, Nadeem and Nishiten Shah (2006) 'Misunderstanding Metaethics: Korsgaard's Rejection of Realism', in Shafer-Landau, Russ (ed.), *Oxford Studies in Metaethics, Vol. 1*, Oxford; New York: Clarendon Press.
- Korsgaard, Christine (2003) 'Realism and Constructivism in Moral Philosophy', *Journal of Philosophical Research, APA Centennial Supplement* 99-122.
- Lynch, Michael (2009) *Truth as One and Many*, Oxford; New York: Oxford University Press.

- Pritchard, Duncan (2005) *Epistemic Luck*, Oxford; New York: Oxford University Press.
- Price, Huw. (2004) 'Naturalism without representationalism', in David Macarthur and Mario de Caro (eds.), *Naturalism in Question*, Harvard University Press, pp. 71—88.
- Quine, W.V. (1970) *Philosophy of Logic*, Englewood Cliffs: Prentice Hall.
- Rawls, John (1980) 'Kantian Constructivism in Moral Theory', *Journal of Philosophy* 77: 515-572.
- 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.
- Schiffer, Stephen (1972) *Meaning*, Oxford: Oxford University Press.
- Schroeder, Mark (2008) *Being For*, Oxford; New York: Oxford University Press.
- Sellars, Wilfrid (1956) 'Empiricism and the Philosophy of Mind', in Scriven, Feigl & (ed.), *Minnesota Studies in the Philosophy of Science*, Minneapolis, MN: University of Minnesota Press.
- Smith, Michael (1994) *The Moral Problem*, Oxford: Blackwell.
- Street, Sharon (2008) 'Constructivism About Reasons', in Shafer-Landau (ed.), *Oxford Studies in Metaethics, Vol. 3*, Oxford; New York: Oxford University Press.
- Stoljar, Daniel (1993) 'Emotivism and Truth Conditions', *Philosophical Studies* 70: 81-101.
- Timmons, Mark (1999) *Morality without Foundations: A Defense of Ethical Contextualism*, New York: Oxford University Press.
- Unwin, Nicholas (1999) 'Quasi-Realism, Negation and the Frege-Geach Problem', *Philosophical Quarterly* 49: 337-362.
- Williams, Michael (1992) *Unnatural Doubts*, Princeton, NJ: Princeton University Press.
- Wright, Crispin (1992) *Truth and Objectivity*, Cambridge, Mass.; London: Harvard University Press.