

Durham Research Online

Deposited in DRO:

17 January 2019

Version of attached file:

Accepted Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Cowie, Christopher (2018) 'The similarity hypothesis in metaethics.', in *Methodology and moral philosophy.* , pp. 127-147.

Further information on publisher's website:

<https://www.routledge.com/Methodology-and-Moral-Philosophy/Suikkanen-Kauppinen/p/book/9781138322097>

Publisher's copyright statement:**Additional information:**

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

DRAFT. See Final Version Forthcoming in Suikkanen (ed.) *Methodology and Moral Philosophy*.

The Similarity Hypothesis in Metaethics

In the introduction to his new book *Being Realistic about Reasons* T.M. Scanlon writes:

“Contemporary metaethics differs in two important ways from the metaethics of the 1950’s and 60’s and even the later 1970’s when John Mackie wrote *Ethics: Inventing Right and Wrong*. In that earlier period, discussion in metaethics focused almost entirely on morality.... Today, although morality is still much discussed, a significant part of the debate concerns... reasons for action and even more broadly reasons for belief and other attitudes.” (2014, 1).

Scanlon is capturing a thought that many contemporary metaethicists would both *agree with* as a sociological description of their subject and *endorse*. We should approach metaethics with the big picture in mind: a picture that includes not just morality, but ‘reasons for belief and other attitudes’. Doing so is interesting in its own right as well as, potentially, a useful means of shedding light on morality itself.

There is a sense in which this represents a new ‘metanormative’ approach to, or method in, metaethics. But what results has it actually had? One interesting example is the widespread use of what I shall refer to as ‘The Similarity Hypothesis’. This is the view that practical normativity – what we ought to do – and theoretical normativity – what we ought to believe – are similar in metaphysically important respects. It is a view that is increasingly appealed to in arguments in metaethics and beyond. The aim of this paper is to assess the Similarity Hypothesis. I do so by engaging with one of the most interesting arguments *against* it.

The paper proceeds as follows. I'll begin by saying a bit more about the Similarity Hypothesis. I'll then introduce the challenge to it that I'm interested in. The challenge concerns the contrasting natures of the 'aims' of belief and desire (or intention) respectively. I'll briefly rehearse and set to one side familiar attempt to articulate this argument before engaging at length with Stephen Darwall's articulation. This will be my main focus. I'll claim that despite its promise, Darwall's argument leaves the Similarity Hypothesis untroubled.

1. The Similarity Hypothesis

According to *The Similarity Hypothesis* theoretical and practical normativity are, at the metaphysical level, the same kind of thing. There are probably ways of reading this claim on which it is trivially true and ways of reading it on which it is trivially false. But I intend something like this:

The Similarity Hypothesis: With respect to the big picture metaphysical questions about their nature and metaphysical grounds, theoretical normativity and practical normativity look similar: for example, if non-natural realism is true about one, then it's true about the other, if expressivism is true about one, then it's true about the other, if one is grounded in our attitudes, or the activity of reasoning, then so is the other. And so on.

This view figures prominently in contemporary metaethics, sometimes as an assumption in arguments, sometimes as defended premise, sometimes as a conclusion.

Example 1: Kearns and Star

Consider, for example, Kearns and Star's recent defence of the view that *reasons are evidence* (RA). The first argument that they offer this is the following (2009, pp.219):

1. Epistemic and practical reasons are of a kind.
2. RA provides the only plausible account of reasons according to which (1) is so.
3. Therefore, RA is true (inference to the best explanation).

This argument clearly makes use of the Similarity Hypothesis as a premise in order to establish a conclusion about the metaphysics of theoretical and practical normativity. (The Similarity Hypothesis is, of course, premise 1).

Example 2: Parfit

We find another nice example in Derek Parfit's critical discussion of a view that he calls 'Metaphysical Naturalism'. This is the view that normative facts are natural facts. Parfit's discussion of this view is almost entirely focused on *practical reasons*. In concluding the section, however, Parfit writes:

"If Metaphysical Naturalism were true, we could not have reasons to have particular beliefs. Such epistemic reasons are also irreducibly normative.... So it could not be true that we *ought* to accept Naturalism, nor could we have any reasons to accept this view." (2011, pp.110).

What's really going on here is that Parfit is using the Similarity Hypothesis to reject a naturalistic view practical reasons. He is effectively arguing as follows:

1. If metaphysical naturalism is true about practical reasons, then it is true about epistemic reasons.
2. If it is true about epistemic reasons, then we would not have any reason to believe it to be true.
3. So, metaphysical naturalism is not true about practical reasons.

We may or may not think that this as a very good argument. But that's not what really matters here. What really matters is that it is another instance in which the

first premise - the Similarity Hypothesis – is being used to draw a big picture metaphysical conclusion.

Example 3: Smith

In the most recent articulation of his view, Michael Smith argues for a particular view of reasons for intentional action; roughly, the view that reasons for action are properly analysed as reasons for believing something to be desirable. His argument is the following:

1. The concept of a reason for belief is analyzable in non-normative terms (i.e. in terms of probability and evidence).
2. So, if the concept of a reason for desire or action were unanalyzable then “the concept of a reason would be a ragbag” (2017, pp.102).
3. The concept of a reason is not a ragbag.
4. The concept of a reason for desire or action is analysable.

This is also an argument in which the Similarity Hypothesis figures as a premise – in this case premise 3. Smith is arguing from the fact that the concept of a reasons for belief and a reason for intentional action are, in some important sense, the same kind of thing, to a conclusion about the metaphysics of reasons for action (roughly, analytical reductionism). Although couched in terms of the similarity of *concepts* (rather than properties), this is, I think, nevertheless recognisable as an appeal to a version of the Similarity Hypothesis.

Example 4: Cuneo

Finally, consider Terence Cuneo’s recent use of the Similarity Hypothesis. Cuneo has argued for the more focused view that *moral* normativity and epistemic normativity are, at the metaphysical level, the same kind of thing. And he has used this as part of an argument for moral realism. His ‘master argument’ is as follows (2007, 6):

1. If moral facts do not exist, then epistemic facts do not exist.
2. Epistemic facts exist.
3. So, moral facts exist.
4. If moral facts exist, then moral realism is true.
5. So, moral realism is true.

A version of the Similarity Hypothesis figures here in premise 1. It would be grossly unfair on Cuneo – and others following similar lines of argument, like Matt Bedke Richard Rowland – to say that the argument *assumes* the Similarity Hypothesis.¹ Cuneo (and others) argue for it at length. Nevertheless, its truth is a (the) key move in their argument.

These examples show that the Similarity Hypothesis, in one form or another, figures prominently in contemporary metaethics; and there are more examples out there. This isn't surprising. It is a sensible hypothesis. After all, both practical normativity and theoretical normativity are kinds of normativity.² Furthermore, there are some properties, like honesty or sincerity, that seem to straddle the border between the theoretical and the practical, making a distinction in their metaphysical natures *prima facie* unattractive.³

2. Fitzpatrick's Point

The Similarity Hypothesis can be challenged. I'll not rehearse all of the challenges here. My aim is to look in detail at one interesting challenge that is occasionally mentioned in the literature, but rarely (if ever) filled out. Consider the following remark from William Fitzpatrick, which nicely summarises the basic idea:

“Suppose... that the existence of epistemic normative facts is tied to the fact that belief has a distinct constitutive aim: namely, truth... an ethical antirealist might grant all this while remaining sceptical that there is any similar

¹ Bedke (2010), Rowland (2013).

² See e.g. Stratton-Lake (2002) and Rowland (2013).

³ This figures prominently in Cuneo's (2007) argument. The metaethical significance of 'entanglement' is also to the fore in Putnam (2002).

analogue of a constitutive aim for... desire or intention..." (Fitzpatrick 2009, pp.757)

I'll call this *Fitzpatrick's Point*.⁴ Fitzpatrick's point is comprised of two claims. The first claim concerns the nature of epistemic normative facts. Roughly:

BELIEF: Belief has a constitutive aim - truth or knowledge. This explains or grounds the existence of epistemic normative facts.

The second concerns the nature of intention or desire and their relation to the existence of practical normative facts. Roughly:

NOT DESIRE: Desire or intention doesn't have a constitutive aim that explains the existence of practical normative facts.

In what follows, I'll discuss the case for thinking that both of these claims – and hence, Fitzpatrick's Point - are simultaneously true.

Three brief clarifications before I begin. Firstly, I will for the most part use the slightly clunky expression 'desire or intention' throughout. The views that I'm discussing focus on one or more of these two, but there's no great continuity in the literature. At some points, for simplicity, I'll just talk about desire (and not intention) if that is the best fit with the authors whose work I am discussing. Nothing, I hope, turns on it. Secondly, I will focus – again, following the authors with whom I am concerned - on the view that belief's aim is truth, rather than knowledge.⁵ This too, I hope, is innocuous in the present context.

Thirdly and most importantly the two different arguments in defence of Fitzpatrick's Point that I'll be discussing are based on two different readings of

⁴ Though it is not a point that Fitzpatrick endorses. He raises it as a plausible source of skepticism about unified approaches to the metaphysics of theoretical and practical normativity.

⁵ For some discussion of whether the aim is truth or knowledge see Wedgwood (2002), Engel (2013), Whiting (2013).

'aim' – one teleological, one normative. On the teleological view attitudinal aims are understood in terms of the intentions of those who hold those attitudes. On the normative view, by contrast, attitudinal aims are understood in terms of the normative properties that (thereby) govern those who hold those attitudes. I'll say more about both in due course. Importantly however, how we develop Fitzpatrick's Point depends on which of these two views of 'aim' we work with. I'll begin by saying something briefly (in the next two sections) about how Fitzpatrick's Point can be developed if we work with the teleological view. But my focus will really be on whether we can develop Fitzpatrick's Point if we work with the normative view.

3. The Teleological View

There is an established literature that is most closely, but not solely, associated with David Velleman's work in the 90's and early 2000's that (a) employs the teleological view and (b) does so in a way that is conducive to Fitzpatrick's Point.⁶ Much of this will probably be familiar to readers, but I'll sketch it in the next two sections nonetheless, before moving on to the normative view.

The key points – up front - are the following.

- (i) Belief aims, in the teleological sense, at truth.
- (ii) Because of this, the nature of belief grounds the existence and nature of reasons for belief.
- (iii) You might think that desire/intention aims (in the teleological sense) at goodness or desirability, and so grounds the existence and nature of practical reasons....
- (iv) ...But you'd be wrong. Desire/intention doesn't aim at goodness or desirability. It aims at something less interesting that can't ground the existence of practical reasons.

⁶ Velleman (1992, 1996, 2000). The view changes in Velleman and Shah (2003). A similar view can be found in Millar (2004).

Putting these together we're meant to get:

- (v) There's good account of the grounds of theoretical normativity in terms of the nature of belief, but no analogous good account of the grounds of practical normativity in terms of desire/intention.

And that – i.e. (v) - is pretty much just the view that the Similarity Hypothesis is false. Let's take each of these claims a bit more slowly (objections to follow).

(i): For belief to aim in the teleological sense at truth is for one to count as believing a proposition just in case one intends to accept it iff it is true. There are several reasons for thinking that this is a good view. One is that it allows for an attractive differentiation between belief and other closely related propositional attitudes such as imagining or supposing. These are attitudes that plausibly characterised as instances of *accepting* a proposition. But they aren't beliefs. And one obvious way of thinking about what differentiates them from beliefs is that when you imagine or suppose a proposition to be true, you needn't be accepting it with the intention of doing so iff it is true. You could be accepting it with the intention of its being, for example, useful, or interesting to do so. Not so with belief. When you believe a proposition you accept it with the intention of doing so iff it is true.

(ii): The fact that belief has this aim grounds the existence and nature of reasons for belief. Specifically, it explains why evidence for a proposition always constitutes a reason to believe that proposition. It does so because believing on the basis of what your evidence supports is the best *means* of satisfying your intention; the intention that you must have insofar as you count as believing a proposition at all. Reasons are generated in this means-ends way; they're explained by being the means to the end of fulfilling the intention that is constitutive of believing.

(iii): Now it is sometimes claimed that much as belief aims at truth, so desire or intention aims at the good or desirable; a view that is, for better or worse,

associated with Aristotle and more recently with Anscombe and Davidson.⁷ This view is sensible on a first pass; it is plausible that if you intend to act in some way, then you must always take acting in that way to be good or desirable *in some sense*. If this were the case, then you might think that the nature of desire/intention could ground the existence and nature of practical reasons. You would have practical reason to take the means to the end of acting in good or desirable ways. But, if the Velleman-inspired literature is correct - you would be wrong to think this. One can desire or intend to act in ways that one doesn't - or wouldn't, on reflection - take to be good in any way. These are *perverse* desires or intentions, and they are probably more common than many of us would like to admit. They are still desires or intentions, however, and we can still perform full-blooded actions on their basis. So desire/intention doesn't aim at the good or desirable.⁸

(iv): If desire/intention doesn't aim at the good or desirable, what does it aim at? That's contested. But the answer is going to be too bland to ground the existence of practical reasons. To see this, note that whatever the aim is, it's got to be something that all desires/intentions have in common, including desires/intentions for what's best, for what's merely somewhat good, and for what isn't good at all (i.e. 'perverse' cases). And whatever fills *that* role is going to be something fairly mundane. Candidates include: that the resulting action have some point or purpose, and that the resulting action be under the control of the actor. Now whilst any one of these aims might ground *some* practical reasons - for example, reasons not to act in ways that are unsatisfiable - it is very hard to see how it could ground *all* practical reasons:⁹

“Theses about the constitutive aim of intentional action are about what is intrinsic to all intentional action. An action may achieve the constitutive aim while being subject to criticism in all sorts of ways that are not explicable just

⁷ Anscombe (2000), Davidson (1980).

⁸ Velleman (Ibid).

⁹ See Millar 2004, Velleman (Ibid). This is not uncontroversial. It has been argued that the nature - if not the aim, exactly - of desire or intention can ground practical reasons quite generally, e.g. Korsgaard (2009).

in terms of the constitutive aim. An action may achieve the constitutive aim while being foolish, or morally wrong, or perverse, or thoughtless.” (Millar, 2002, pp.68).

(v): Putting these points together, it looks as though we have a good account of the grounds of theoretical normativity but no analogous good account of the grounds of practical normativity. Specifically, theoretical normativity is grounded in the nature of belief, whereas practical normativity is not grounded in the nature of desire or intention. It remains to be seen how practical normativity *is* grounded. But we know that the story will have to be very different from the story of how theoretical normativity is grounded. This clearly threatens the Similarity Hypothesis.

4. The Problems for the Teleological View

This simple argument doesn't work though, or at least it doesn't obviously work. So a defender of the Similarity Hypothesis needn't be too worried by it. There are problems with all of (i)-(iv). But just think about the main problems with (i) and (ii) to start.

As regards (i): It is not at all clear that the proposed account of belief – as aiming at truth in the teleological sense – is correct. There are lots of reasons for this. One reason is that only a very small subset of our beliefs – roughly, those that are a result of conscious deliberation - seem to be accompanied by an *intention* of the kind that the teleologist needs. Ordinary beliefs formed on the basis of perception, say, don't seem to be like this. Defenders of the teleological account have responded by claiming that sub-personal belief-forming systems can in some relevant sense be governed by, or overseen by, our intentions with respect to them, or that these systems can be ascribed 'aims' in selective senses.¹⁰ But this is an unhappy account in various respects. It's a real push to claim that we have higher-order intentions with respect to our sub-personal belief-forming systems. And if we make recourse to understanding the aims of sub-personal

¹⁰ Velleman (2000), Steglich-Petersen (2006).

systems in terms of their selective functions, then it's hard to see how these could ground normative reasons (as they would need to for BELIEF to be true).¹¹

A second problem is that the teleological characterisation of belief is too strong in an obvious respect. The teleological account – at least as I've presented it – rests on the supposed datum that states that are sensitive to our wishes as well as our evidence are *not beliefs*. But this isn't an obvious pretheoretical datum at all. There are lots of states that we ordinary call beliefs – albeit bad ones – that are sensitive to our wishes in some way. Shah puts this nicely:

“Our discourse is replete with accusations that likes and dislikes influence belief... Recent debates about the alleged connections between IQ and poverty and sociobiological explanations of human sexual behavior are rife with such accusations on both sides, and while individual charges of prejudice might be disputed, the legitimacy of the general form of such criticisms is taken for granted by all sides.” (Shah 2003, pp.16)

Now it would be possible to modify the teleological account to take care of this: in order to count as believing a proposition one's intention with respect to it must be merely *somewhat* directed at truth. But this modification runs into problems elsewhere. It is not longer obvious how we could use the resulting account of belief to differentiate beliefs from closely related attitudes (such as imagination or supposition) or to explain the impossibility of believing at will.

As regards (ii): On the view discussed above, reasons (for belief) are generated by intentions via roughly the following principle: if φ -ing is the best means to Ψ -ing, and if you intend to Ψ , then you have a reason to φ . This is a highly contestable principle however; one that would require substantial further defence.¹²

¹¹ Papineau (2013).

¹² Enoch (2011).

So where does this leave us? It's highly questionable whether the teleological view can establish BELIEF and *a fortiori* highly questionable whether it could support Fitzpatrick's Point.

5. The Normative View

The alternative normative reading of 'aim' is often thought to fare better as an account of the sense in which belief aims at truth. The basic idea is usually put like this:

N_B : A belief is correct if and only if it is true.

Correctness here is a normative property. It denotes a normative standard for our responses (e.g. beliefs, desires, intentions) that holds just in virtue of one's being engaged in the kind of enterprise that one is engaged in. There is some debate about how exactly it should be understood – whether in terms of reasons, oughts, evaluative properties (e.g. goodness) or 'fittingness'. But I'll bypass these issues here.¹³

N_B is an attractive way of thinking about the sense in which belief aims at truth. For one thing it promises to avoid or resolve some of the problems with the teleological reading. For example, N_B is not committed to the view that we or our sub-personal systems must possess intentions with respect to belief-formation. N_B is also not committed to the troublesome view that beliefs are always influenced solely by truth-related considerations, rather than by our desires. A defender of N_B can, however claim that when our beliefs are formed on the basis of our desires rather than our evidence, they're not as they *ought* to be.¹⁴

¹³ For discussion see e.g. McHugh and Whiting (2014), Greenberg and Cowie (2016).

¹⁴ Defenders of N_B have also claimed, plausibly, that their view can explain a range of further features of belief, including transparency in deliberation and Moore-paradoxical sentences.

A further attraction of N_B is that it is well placed to explain the existence and nature of epistemic reasons for belief. For a defender of N_B , reasons to believe the truth and disbelieve falsehoods will fall directly out of the nature of belief. The mere fact that belief is the kind of thing that it is explains this. Reasons to believe what one's evidence supports and disbelieve evidentially unsupported propositions will be explained as subsidiary norms of N_B ; roughly, as norms that one would follow in order to conform to N_B given one's epistemic limitations.¹⁵

For these reasons, there is a lot to be said for the view that BELIEF is true in virtue of N_B . There are also, I should add, serious objections to the view. I'll return to this in conclusion. But suppose in any case that we do accept N_B and, on this basis, BELIEF (as many seem to). My interest is in whether this could be the basis for an argument for Fitzpatrick's Point. In order to do so, we would need to see what the normative view of 'aim' implies regarding the aim of desire/intention.

Much the most plausible and widely accepted view is that it would imply that the aim of desire or intention is the good or desirable after all. Or:

N_A : An action (or intention or desire) is correct iff it would be good or desirable.¹⁶

This is plausible because, whilst it might not be that you do always desire or intend the good or desirable, presumably you *ought* to. So if we shift from the teleological account of 'aim' to the normative 'aim', then it becomes plausible that the aim of desire is goodness or desirability after all.¹⁷

¹⁵ For the full case for N_B see e.g. Wedgwood (2002, 2007), Shah (2003). For helpful discussion see Engel (2013), McHugh and Whiting (2014), Greenberg (2017).

¹⁶ See e.g. Shah (2008), Evans and Shah (2012). Wedgwood (2003, 2007) makes much the same point but puts it in terms of 'choice' and 'choiceworthiness'.

¹⁷ It is also sometimes argued that N_A is supported by a certain kind of 'transparency' in the nature of practical deliberation illustrated by the toxin puzzle (Shah 2008).

All of this would seem to undermine NOT DESIRE, and hence Fitzpatrick's Point. If N_A were true, then the nature of intention or desire would seem to have just as good a claim to ground practical reasons as the nature of belief would to ground epistemic reasons. You'd have practical reasons to act in desirable ways, much as with N_B you'd have reasons to believe the truth and avoid falsity (i.e. epistemic reasons).

6. Darwall

Enter Darwall.¹⁸ Darwall is happy to accept that the aim of desire is the good or desirable. But he doesn't think that this does any work in grounding practical reasons. And he thinks that this marks a difference between belief and desire. The aim of belief *does* do some work. In other words, Darwall seems to provide the means to defend Fitzpatrick's Point in a way that looks consistent with the normative reading of 'aim' after all.

Darwall's argument is based on a distinction – borrowed from Velleman – between two very different ways of specifying the aim of a practice. The aim can be specified in either *formal* or *substantial* terms. The distinction is best illustrated by example. There's a fairly uninformative sense in which the aim of chess – like many other games and sports – is to win. That is its *formal aim*; winning. The substantial aim is what that formal aim consists in. In the case of chess, it is check-mating the opposing King. Different games will, of course, have different substantial aims. In Snap it is to match more pairs of similar cards than your opponent, in mixed martial arts it is to incapacitate your opponent by knockout or submission, and so on.

Now think about the claims that belief is the aim of truth and that desirability is the aim of desire respectively. These are disanalogous in that the former is about a substantial aim whereas the latter is about a merely formal aim. The formal aim of belief is presumably belief-worthiness, and the substantial aim of desire is

¹⁸ Darwall (2003).

whatever desirability consists in (which could be any number of things – more on this below). This disanalogy is explicitly appealed to by Darwall:

“Owing to the nature of belief... objective truth and probability bear on what we have reason to believe... Still the concepts of truth and objective probability differ from the normative concept of what one ought to believe... whereas believing as we ought is belief’s “formal aim” truth is belief’s “substantive aim”.

The situation is different in the practical case... Evaluative and practical attitudes do, like belief, have their respective “formal aims”. The formal aim of desire is the desirable... of choice, the choiceworthy... But unlike belief, no evaluative or practical attitude has a “substantive aim” ... (pp.484)

Furthermore – and this is the crucial point – Darwall seems to take this disanalogy to be of metaphysical significance. It supports the view that theoretical reason and practical reason are, at the big-picture metaphysical level, very different things. Darwall writes:

It follows that, unlike theoretical reason, practical reason is responsible to no external goal or standard which logically closes sensible deliberation. Reasons for acting can be grounded nowhere but within norms of free practical reason itself.” (pp.484-5).

This is a lot going on here. But Darwall certainly appears to be presenting an argument for Fitzpatrick’s point – and so, against the Similarity Hypothesis - based on the contrast between the substantial nature of belief’s aim on one hand and the merely formal nature of desire’s aim on the other. But what, exactly, is the argument, what is it meant to show, and how plausible is it? The argument, as I see it, is comprised of the following core claims:

- (1) Because belief has a substantial constitutive aim, theoretical reason is responsible to 'an external standard (truth) that logically closes sensible deliberation'.
- (2) Because desire doesn't have a substantial constitutive aim, practical reason is *not* responsible to such a standard.
- (3) Given that – from (2) - practical reason is *not* responsible to such a standard, it must be grounded in 'the norms of free practical reason'.
- (4) By contrast, given that – from (1) - theoretical reason *is* responsible to such a standard, it is not grounded in 'norms of free theoretical reason'.

I hope that this is an accurate rendering of Darwall's argument. Given my own aims in this paper, I am interested to see whether it could be used to support the further claims:

- (5) From (3) and (4), the grounds of practical reason and theoretical reason are of fundamentally different kinds.
- (6) So Fitzpatrick's Point is right and the Similarity Hypothesis is wrong.

In the remainder of this paper I will try to find out whether this is a good argument. Before doing this it is worth noting that – although I'll focus on Darwall's argument - this *kind* of argument isn't unique to Darwall. I think that we find a very similar thought in Hallvard Lillehammer's brief discussion of the prospects for drawing an analogy between 'epistemology' and 'ethics'. Sceptical of this analogy, Lillehammer writes:

"While arguments in epistemology often have the pre-theoretical appearance of being at least partly controlled by some conception of an external aim or function usually described in terms of concepts such as 'truth' or 'knowledge',

there does not seem to be an obviously analogous aim or function external to arguments in ethics... ” (2007, pp.170).

So both Darwall and Lillehammer seem to have – at least – sympathy with the same kind of reading of Fitzpatrick’s Point (indeed Fitzpatrick makes his point in the specific context of noting the above quotation from Lillehammer). It is, then, worth taking seriously. Furthermore it is an argument that I have a lot of sympathy with on a first reading. But does it work? I’ll argue that it doesn’t.

7. Assessing Darwall’s Argument: Premise 1

Begin with:

- (1) Because belief has a substantial constitutive aim, theoretical reason is responsible to ‘an external standard (truth) that logically closes sensible deliberation’.

This premise states that theoretical reasoning has an external standard, and that it has it in virtue of the nature of belief. But what is an ‘external standard’ for theoretical reasoning? This requires some interpretation. A good place to begin is, I think, by thinking about what an *internal standard* for theoretical reasoning would be. By a ‘standard’, I’m just going to understand an aim or goal for that practice, which in turn I’m going to understand in normative terms. So the standard for a practice is its correctness condition. ‘Internal’ is harder. There is no precise, established usage here so a bit of interpretation is required. But I’m going to suggest that for a standard to be ‘internal’ to theoretical reasoning would be for that standard to be, in some important way, dependent on doing some (actual or hypothetical) theoretical reasoning. And there are two obvious ways in which the standard of theoretical reasoning could be so dependent. One is epistemological. It is that doing some theoretical reasoning is necessary to *identify* or *work out* what the standard of theoretical reasoning is. The other reading is metaphysical. It is that something’s status as being the standard of theoretical reasoning is metaphysically dependent on some theoretical

reasoning that someone might (actually or hypothetically) do. For example, perhaps being the standard of theoretical reasoning just is being the view of what you ought to believe that you would arrive at by thinking about what you ought to do.

So what about an *external* standard, and the claim that theoretical reasoning is responsible to such a standard? This can be read in either epistemological or metaphysical terms, mirroring the epistemological and metaphysical readings of 'internal standard' above. On the epistemological reading, it's that you *don't* need to do any theoretical reasoning to work out what the standard of theoretical reasoning is. On the metaphysical reading, it's that something's status as being the standard of theoretical reasoning *isn't* metaphysically dependent on some theoretical reasoning that someone might (actually or hypothetically) do. So we can interpret premise 1 in either of the following two ways, firstly epistemologically:

1.2_E Because belief has the substantial constitutive aim of truth, you don't have to do any theoretical reasoning *to work what* or *identify* what you ought to believe (i.e. truth and not falsehood).

Secondly, metaphysically:

1.2_M Because belief has the substantial constitutive aim of truth, the facts about what you ought to believe aren't such in virtue of your doing any theoretical reasoning.

Which of these readings is Darwall going for? Which *should* he go for? Either would seem very plausible.

Begin with the metaphysical reading. Unless a pragmatist view of truth – i.e. truth is acceptability to ideal enquirers – is correct, then the truth of propositions isn't, in general, metaphysically dependent on one's doing any theoretical reasoning. Darwall doesn't assume pragmatism about truth, nor is it

obvious that he or anyone else should. The idea behind the epistemological reading in contrast is that because belief aims constitutively at truth, working out that a proposition is true suffices for working out that you ought to believe it. Given the nature of belief, if you work out that a proposition is true, you don't need to do any further enquiry to work out that you ought to believe it.

One reason to think that Darwall is really going for the epistemological reading is that it best explains his additional claim about *deliberation*. The fact that belief aims at truth, he claims, 'closes sensible deliberation'. This is, to my mind, most naturally read as supporting the epistemological reading. To say that the fact that belief aims at truth 'closes sensible deliberation' is pretty much just to say that a sensible deliberator will, on having deliberated and concluded that some proposition is true, end her deliberations there; she'll believe that proposition. She won't then go on to ask "but should I believe it?".¹⁹

So the essence of 1, as I understand it, is this: because belief aims at truth, when you've worked out that p is true you don't need to do any additional thinking

¹⁹ There is a slight complication here. This argument will arguably only work if the fact that belief aims constitutively at truth is part of, or is entailed by, the *concept* of belief. To see why this is important, distinguish between two ways in which belief's aim can be constitutive. Firstly, its aim can be constitutive of what it is to believe something but not constitutive of the concept of belief. When this is the case, the aim is *metaphysically constitutive* (think water = H₂O, but for constitutive aims). Secondly, its aim can be constitutive of both what it is to believe something and of the concept of belief. When this is the case, the aim is *conceptually constitutive* (think bachelors = unmarried men, but for constitutive aims). The conceptually constitutive reading would arguably be necessary if the aim of belief is to 'close sensible deliberation'. That's because deliberation won't be sensibly closed by finding out that a proposition is true unless the deliberator knows that a belief ought to be true (and not false). And whilst it is plausible to ascribe conceptual competence to a sensible deliberator it isn't plausible to ascribe competence about the (metaphysical and non-conceptual) nature of things. I don't think, though, that this slightly complication seriously undermines the basic case for thinking that Darwall is really going for the epistemological reading. The point about 'closing deliberation' suffices to establish that. Furthermore, the most prominent and persuasive arguments for N_B are, in fact, arguments for the conceptually constitutive reading and not merely the metaphysically constitutive reading. See McHugh and Whiting (2014).

about what you ought to believe to work out that you ought to believe p. You can just stop there. This seems clear and plausible.

8. Assessing Darwall's Argument: Premise 2

With this in mind, let's now turn to:

(2) Because desire doesn't have a substantial constitutive aim, practical reason is *not* responsible to 'an external standard that logically closes sensible deliberation'.

This is structurally similar to (1). And much as with (1), the main issue here is that we need to separate out the metaphysical and epistemological readings. On the *metaphysical reading* this premise should be read as making a claim about the nature of the formal aim – desirability – itself. It should be read as claiming that what it is to be desirable just is to be the result of some doing some practical reasoning. This view falls naturally within the traditional *constructivist* family of views in metaethics.²⁰ We can understand it roughly as follows:

2_M Because desire's constitutive aim isn't substantial, the facts about what you ought to do are such in virtue of your doing some practical reasoning reasoning.

On the *epistemological reading* by contrast this premise should be read as making a much more modest claim: a claim about how we *work out* or *identify* what the substantial nature of the formal aim (desirability) is. It should be read as claiming that in order to work out the substantial nature of desirability, you need to engage in practical reasoning. This is compatible with the great majority of views in traditional metaethics – constructivist, realist, and otherwise. We can understand it roughly as follows:

²⁰ Street (2010).

2_E Because desire's constitutive aim isn't substantial, you have to do some practical reasoning *to work out* or *identify* what you ought to do.

So which of these readings does Darwall go for? Which should he go for? There are two reasons to think that he should go for the weaker, epistemological reading.

One reason is that it is a good fit with the epistemological reading of (1), which, recall, explained Darwall's claim about 'closing deliberation'. Far more importantly however, Darwall should go with 2_E because it is not at all clear that Darwall would be *entitled* to the stronger metaphysical reading, 2_M. He wouldn't be entitled to it because the fact that desire's constitutive aim isn't substantial *simply doesn't entail* that the nature of desirability is metaphysically dependent on the outcome of practical reasoning. It simply doesn't entail, in other words, the truth of constructivism. To see this, it is useful to think about traditional 'realist' views in moral philosophy. According to these views – unlike constructivist views – what one ought to do is *not* metaphysically dependent on the result of doing some practical reasoning. Rather it is dependent on attitude or reasoning independent normative facts. But these realists could presumably quite happily accept that desirability is the constitutive, formal aim of desire. Yet if 2_M is right, then these realists would thereby be displaying some kind of incoherence. This seems highly unlikely. At the very least, it would require *a lot* more argument. Much more plausible is the weaker epistemological reading, 2_E. On this reading, it is necessary to do some practical reasoning to *work out* what one ought to do. This is straightforwardly compatible with both constructivism and realism and is, in fact, exactly what both of these schools claim. So I think we should certainly read 2 as 2_E.

9. Assessing Darwall's Argument: Premise 3

With this in mind, we can draw some conclusions about the third premise of Darwall's argument, and indeed the nature of his argument more generally. According to the third premise:

(3) Given that – from (2) - practical reason is *not* responsible to an external standard, it must be grounded in ‘the norms of free practical reason’.

Is this a metaphysical claim? As a metaphysical claim it would be a claim about the metaphysical grounds of facts about what one ought to do. It would state that the *metaphysical grounds* of facts about what one ought to do is ‘the norms of free practical reason’ (obviously, we’d need to hear more about what this is). Or is it an epistemological claim? As an epistemological claim it would be a claim about *how to find out* what one ought to do? It would state that we must do so by using ‘the norms of free practical reason’.

To me it sounds very much like Darwall is intending it as a *metaphysical* claim (perhaps an epistemological claim too, but certainly a metaphysical claim). It sounds like a claim about the metaphysical grounds of facts about what one ought to do. More specifically, it sounds like the claim that one ought to do is metaphysically dependent on the outcome of one’s practical reasoning on the model of the traditional constructivist school. Indeed, I think that further features of the paper (and Darwall’s own work independently) speak to this intention.²¹

Suppose that I am right about this. It would be a problem for Darwall. It would be a problem because it wouldn’t fit with the conclusions that I drew about (2). In discussing (2) I argued that the sense in which Darwall could feasibly claim that desire isn’t, by its nature, responsible to an external standard is an *epistemological* sense. It is that we need to engage in practical reasoning to find

²¹ To see this note, very briefly, that the conclusion that Darwall goes on to draw in the later section of the paper is that (some) practical reasons, unlike theoretical reasons are *agent-relative* and not merely *agent-neutral*. That he draws this conclusion is, I think, an indicator that he is reading (3) in the *metaphysical* sense. Whether *there are* any agent-relative reasons - and not merely how we could know about them if there were any - is clearly a metaphysical and not merely an epistemological matter (I won’t detour to the details of the argument here). See also Darwall (1992) for a discussion of the kind of constructivist view that he favours.

out what we ought to do. But I'm now claiming that the conclusion that he is drawing in (3) is a *metaphysical* conclusion. It is, effectively, the conclusion that practical reasoning is metaphysically grounded in 'the norms of free practical reason'. So it looks to me as though there has been a slide from a plausible epistemological claim, to a much more controversial metaphysical claim; a metaphysical claim that, as I argued in discussing (2) – *isn't* obviously supported by the fact that the constitutive aim of desire is merely formal. So if I've read it right, Darwall's argument is actually *invalid*.

10.A Preliminary Conclusion

I may or may not have interpreted Darwall correctly here. I hope that I have, but in any case, from the perspective of my own interests in this paper, the discussion of Darwall's argument above is highly significant.

My own interest, recall, is with whether the Similarity Hypothesis is true and more specifically with whether Fitzpatrick's Point represents an important challenge to it. The reason for discussing Darwall's argument is that it looked like the best shot at making good on Fitzpatrick's Point. If the argument succeeded, then it certainly would promise to do that. I have argued, however, that Darwall's argument doesn't vindicate Fitzpatrick's Point. In order for it to do so, Darwall would need to draw a conclusion about the metaphysics of practical normativity from premises that don't license such conclusions. Perhaps, however, the argument can still be saved. One response would be to argue that I was too quick in my discussion of (2). Perhaps we can make a case for the metaphysical reading of (2) after all. A second response is see if we can use the more modest epistemological reading of (2) to undermine the Similarity Hypothesis in another way. I'll briefly discuss each in turn.

11. Reassessing (2): Could it be a Metaphysical Claim?

I claimed above that (2) should be read as an epistemological claim rather than a metaphysical claim. That was the key to explaining why (3) must be interpreted

epistemologically rather than metaphysically. But perhaps I am wrong about this. Perhaps (2) can be read as a metaphysical claim after all. If it can, then it will license the metaphysical reading of (3) and the subsequent argument against the Similarity Hypothesis.

What are the prospect for this. Recall the two readings. According to the epistemological reading:

1.2_E Because desire's constitutive aim isn't substantial, you have to do some theoretical reasoning *to work out* or *identify* what you ought to do.

And according to the metaphysical reading:

2_M Because desire's constitutive aim isn't substantial, the facts about what you ought to do are such in virtue of your doing some practical reasoning.

Let's not contest the epistemological reading. Let's just ask whether there might be something more to be said for the metaphysical reading. My argument against it consisted of the following claims. Firstly, the fact that desire's constitutive aim isn't substantial doesn't straightforwardly entail that the facts about what you ought to do are such in virtue of your doing some practical reasoning. Secondly, there are plenty of metaethicists who would disagree with this inference in any case; notably, realists. They may well agree that desire's constitutive aim is merely formal, and yet *not* think that the facts about what you ought to do are such in virtue of your doing some practical reasoning.

How might Darwall respond? He would need to show that the above mentioned philosophers are mistaken – perhaps they have made an assumption that they are not in fact entitled to. And to show this he would need to support a missing premise along the following lines:

(C) If the constitutive aim of a practice is merely formal, then the substantial aim of that practice is such in virtue of an appropriate application of the norms of that practice.

If it were possible to provide some general support for this, then he could hope to apply it to the case of desire and substantiate the metaphysical reading of (2) and (3). I don't have any general argument against this principle. But unfortunately, I don't have any positive argument for it either. And the burden of proof is presumably on someone who would argue for it. I leave it as a task for another day (or another philosopher).

12. An Epistemological Threat to the Similarity Hypothesis?

An alternative approach would be to settle for the epistemological reading of (2) and make do with an epistemological reading of (3) as well. An epistemological reading of (3) would just state that the way that we work out what we ought to do is by using the 'norms of free practical reason'. I take it that this just means that the way that we work out what we ought to do is by actually doing some practical reasoning. Now I noted above that this would be a disappointment in two respects. Firstly it would be a disappointment *for Darwall* in that it wouldn't capture his intentions. Secondly, it would be a disappointment *for me* in that it would show that my attempt to substantiate Fitzpatrick's Point by hijacking Darwall's argument hasn't worked.

I don't want to speculate any more on Darwall's intentions. But it is worth saying something more about the second disappointment. As I set it up, the Similarity Hypothesis concerns big picture *metaphysical* differences between theoretical and practical normativity. And I have argued that Darwall's argument isn't going to undermine *that*. But it could undermine a broader Similarity Hypothesis – a hypothesis according to which the similarity between theoretical normativity and practical normativity is *either* metaphysical *or* epistemological. That is to say, it could undermine the following kind of view:

The Expanded Similarity Hypothesis: With respect to the big picture metaphysical *and epistemological* questions about their nature, metaphysical grounds, *and the means of our knowledge about them*, theoretical normativity and practical normativity look similar: for example, if non-natural realism is true about one, then it's true about the other, if expressivism is true about one, then it's true about the other, if one is grounded in our attitudes, then so is the other, *if we know about one in some kind of way, then we know about the other in the same kind of way.* And so on.

In fact, it is very plausible that Darwall's argument does undermine this kind of view. To deny that it does, one would seemingly have to either deny that belief aims constitutively (normatively interpreted) at truth, or show that desire's constitutive aim is not merely formal. And whilst one certainly *could* do this there's some kind of prima facie burden on one who would do so.

The question, however, is whether the Expanded Similarity Hypothesis is at all interesting; whether it is the kind of view that anyone is likely to hold or use in an argument. This is less clear. The examples that I gave of the use of the Similarity Hypothesis at the beginning of this paper were all metaphysical, not epistemological. So they certainly don't seem to rely on or argue for the Expanded Similarity Hypothesis. So perhaps it isn't terribly interesting.

Perhaps, though, there would be *some* mileage to be gained from showing the Extended Similarity Hypothesis to be false. Firstly, it might be interesting just insofar as epistemological differences between theoretical normativity and practical normativity are themselves interesting. You might, for example, be interested in comparing our knowledge of what we ought to do with our knowledge of what we ought to believe, or in thinking about the relative scope of disagreement about what we ought to do and what we ought to think. Consider, for example, Guy Fletcher's suspicion that:

“Whilst cultures differ in what kinds of things they take to be evidence, this does not equate to a body of epistemic norms that is both internally coherent

and fundamentally unlike ours, in the way that ethical norms appear to be. If this is so then, for whatever similarities could be found between ethics and epistemology, there is a countervailing dissimilarity to be considered....” (2009, pp.368).

Fletcher doesn't provide evidential support for this claim, but it is interesting to think that it is plausible that the disanalogy that he anticipates would actually be *predicted* by the nature of the contrasting natures of the attitudes themselves as discussed above.

And there is a second, more ambitious respect in which the denial of the Extended Similarity Hypothesis might be of interest. It might be of interest because the *epistemology* of the some realm of the normative could itself figure in an abductive argument about the *metaphysics* of that realm, as is often the case in 'arguments from disagreement' and 'arguments from epistemic access'. It could thereby be *indirectly* relevant to assessing the Similarity Hypothesis itself. This kind of argument – from epistemology to metaphysics - is not unfamiliar in contemporary metaethics; think, for example, of Mackie's arguments from epistemic access and relativity to his error theoretic conclusion, or more recently of Sharon Street's argument from the epistemological problems for realists (really, a version of the access problem) to the truth of constructivism.²² But of course establishing this would require substantial further work.²³

13. Conclusion

I've tried to give Darwall's argument – and by extension, Fitzpatrick's Point – a good run for its money. Ultimately, I'm not sure that we do end up with a reason for rejecting the Similarity Hypothesis, though I have claimed that we might well have a case for rejecting the Expanded Similarity Hypothesis.

²² Mackie (1977), Street (2006).

²³ For an unsuccessful attempt to establish it in the context of disagreement arguments see Cowie (2013).

Bibliography

- Anscombe, G. E. M. (2000). *Intention*. Cambridge, MA: Harvard University Press.
- Bedke, M. (2010). Might All Normativity be Queer?. *Australasian Journal of Philosophy*, 88(1), pp.41-58.
- Cowie, C. (2013). Epistemic Disagreement and Practical Disagreement. *Erkenntnis*, 79(1), pp.191-209.
- Cuneo, T. (2007). *The Normative Web*. New York: OUP.
- Darwall, S. (1992). Internalism and Agency. *Philosophical Perspectives*, (6), pp. 155-174.
- Darwall, S. (2003). Moore, Normativity, and Intrinsic Value. *Ethics*, 113(3), pp.468-489.
- Davidson, D. (1980). How Is Weakness of the Will Possible?. In: D. Davidson, ed., *Essays on Actions and Events*. New York: Oxford University Press, pp.21-43.
- Engel, P. (2013). In Defence of Normativism about the Aim of Belief. In: T. Chan ed., *The Aim of Belief*. New York: OUP.
- Enoch, D. (2011). Schmagency Revisited. In M. Brady ed., *New Waves in Metaethics*. London, Palgrave-Macmillan.
- Evans, M. and Shah, N. (2012). Mental Agency and Metaethics. In R. Shafer-Landau, ed., *Oxford Studies in Metaethics*, 7, pp.80-109.
- FitzPatrick, W. (2009). Recent Work on Ethical Realism. *Analysis*, 69(4), pp.746-760.
- Fletcher, G. (2009). "Uneasy Companions. *Ratio*, 22(3), pp.359-368.
- Greenberg, A. (2017). The Truth Norm Account of Justification. PhD Dissertation. University of Cambridge.
- Greenberg, A. and Cowie C. (2016). Is the Norm on Belief Evaluative? A Response to McHugh. *Pacific Philosophical Quarterly*. DOI: 10.1111/papq.12158.
- Kearns, S., and Star, D. (2009). Reasons as Evidence. In: R. Shafer-Landau, ed., *Oxford Studies in Metaethics*, 4, Oxford: Oxford University Press., pp. 215-242.
- Korsgaard, C. (2009). *Self Constitution: Agency, Identity and Integrity*. New York: Oxford University Press.

- Lillehammer H. (2007). *Companions in Guilt: Arguments for Ethical Objectivity*. London: Palgrave Macmillan.
- Mackie, J.L. (1977). *Ethics: Inventing Right and Wrong*. New York: Penguin.
- McHugh, C. and Whiting, D. (2014). The Normativity of Belief. *Analysis*, 74(4), pp.698-713.
- Millar, A. (2004). *Understanding People: Normativity and Rationalizing Explanation*. Oxford: OUP.
- Papineau, David. (2013). There are No Norms of Belief. In: T. Chan (Ed.) *The Aim of Belief*. New York: OUP.
- Parfit D. (2011). *On What Matters*. New York: OUP.
- Putnam, H. (2002). *The Collapse of the Fact/Value Dichotomy and Other Essays*. Harvard: Harvard University Press.
- Rowland, R. (2013). Moral Error Theory and the Argument from Epistemic Reasons. *Journal of Ethics and Social Philosophy*, 7(1), pp. 1-24.
- Scanlon, T.M. (2014). *Being Realistic About Reasons*. Oxford: OUP.
- Shah, N. (2003). How Truth Governs Belief. *Philosophical Review*, 112(4), pp. 447-482.
- Shah, N. (2008). How Action Governs Intention. *Philosophers' Imprint*, 8, pp. 1-19.
- Shah, N. and Velleman, D. (2005). Doxastic Deliberation. *Philosophical Review*, 114(4), pp. 497-534.
- Smith, M. (2016). Parfit's Mistaken Metaethics. In: P. Singer, ed., *Does Anything Really Matter? Essays on Parfit on Objectivity*. New York: Oxford University Press., pp. 99-119.
- Steglich-Petersen, A. (2006). No Norm Needed: On the Aim of Belief. *Philosophical Quarterly*, 56(225), pp. 499-516.
- Stratton-Lake, P. (2002). Introduction. In: P. Stratton-Lake, ed., *W. D. Ross. The Right and the Good*. Oxford: Clarendon Press.
- Street, S. (2006). A Darwinian Dilemma For Realist Theories Of Value. *Philosophical Studies*, 127, pp. 109-166.
- Street, S. (2010). What is Constructivism in Ethics and Metaethics?. *Philosophy Compass*, 5(5), pp.363-384.
- Velleman, D. (1992). The Guise of the Good. *Nous*, 26(1), pp. 3-26.

Velleman, D. (1996). The Possibility of Practical Reason. *Ethics*, 106(4), pp. 694-726.

Velleman, D. (2000). On the Aim of Belief. In: D. Velleman, *The Possibility of Practical Reason*. Oxford: Oxford University Press., pp. 244-281.

Wedgwood, R. (2002). The Aim of Belief. *Philosophical Perspectives*, 16(s), pp. 267-297.

Wedgwood, R. (2003). Choosing Rationally and Choosing Correctly. in S. Stroud & C. Tappolet, eds., *Weakness of Will and Practical Irrationality*. Oxford University Press., pp. 201-229.

Wedgwood, R. (2007). *The Nature of Normativity*. Oxford: Oxford University Press.

Whiting, D. (2013). Nothing But the Truth: On the Norms and Aims of Belief. In T. Chan, ed., *The Aim of Belief*. New York: OUP.