

Are all normative judgements desire-like?

The definitive version of this paper can be found on the webpage for the journal of ethics and social philosophy: <http://jesp.org/>.

Abstract. In this paper I first argue against one attractive formulation of the motivation argument, and against one attractive formulation of non-cognitivism. I do so by example: I suggest that other-regarding normative judgements do not seem to have motivational powers and do not seem to be desires. After defending these two claims, I argue that other views can accommodate the motivational role of normative judgement without facing this objection. For example, desire-as-belief theories do so, since such theories only say that *some* normative judgements constitute desires, not that all such judgements do so. (I also briefly present similar reasoning in favour of the claim that desire-as-belief is superior to non-cognitivism with respect to the Frege-Geach objection.) In short, I argue that if we are seeking a theory that explains the motivational role of normative judgement, some theories are better than others insofar as they do so without committing to the claim that *all* normative judgements play such a motivational role.

If I come to think that I ought to go to Sweden, we might think that this judgement is somewhat appetitive: if I really think this, I must be somewhat inclined to go. But in contrast, if I judge that *you* ought to go to Sweden, it is far less clear that this involves any kind of inclination on my part: I might really think that you ought to go, but need not be at all in favour of your doing so (indeed, perhaps I would much prefer you to shirk your duties and stay). Other-regarding normative judgements seem to be a matter of mere recognition, not inclination. This casts doubt on non-cognitivist views according to which all normative judgements are desire-like. But it fits much better with theories in the vicinity of desire-as-belief, which identify only *some* normative judgements with desires. So I shall argue.

The paper is split into seven sections. Section 1 describes a natural way of formulating non-cognitivism, which I label conativism. Section 2 describes the motivation argument, and presents a version that escapes some standard criticisms of that argument. In section 3, I argue that other-regarding normative judgements present a problem for the motivation argument, and indeed present a problem for conativism itself. Sections 4 and 5 consider two possible replies. Section 6 very briefly describes how the problem relates to the Frege-Geach problem. Section 7 argues that some other theories – such as desire-as-belief – may be able to accommodate the motivational role of normative judgement without falling prey to the same problem.

1. Conativism

People have a variety of views about what's good, bad, right, wrong, justified, and so on. It's helpful to think of these as views about normativity (where this may include but is certainly not exhausted by, moral normativity).¹ I'll follow tradition and stipulatively use the word "judgement" to refer to the state of mind (whatever it is) that such views consist in. With this terminology, we can formulate a theory:

Conativism. All normative judgements are desires.

We can think of conativism as one particular kind of non-cognitivist theory. Non-cognitivists deny that normative judgements are beliefs. Conativism adds to non-cognitivism by also making a claim about what normative judgements are: desires. I take it that conativism represents a central strand of the non-cognitivist tradition. For example, one classic argument for non-cognitivism is the motivation argument, which appeals to the fact that normative judgements motivate us in a way that only desires can (see e.g. Blackburn 1984: 187-9, as well as Smith 1994). If this moves us to accept non-cognitivism, it should move us to accept the conativist kind of non-cognitivism, since it establishes the conclusion that normative judgements are desires, not merely the conclusion that they are not beliefs.

In this paper, I object to conativism. I thereby leave open that there might be other non-cognitivist views which are plausible and which escape the objection I present against conativism. For example, I shall not discuss non-cognitivist views that abandon the motivation argument entirely and treat all normative judgements as states of mind that are neither beliefs nor desires. I shall also not discuss non-cognitivist views that claim that only some normative judgements are desires and that others are some other non-cognitive state of mind. I tend to think that conativism captures an important strand of the non-cognitivist tradition, and that non-cognitivist views other than conativism are likely to lose some of the advantages that non-

¹ Throughout this paper, I wholly ignore epistemic normativity, which raises too many issues to be adequately discussed here. To the extent that a conativist analysis of epistemic normative judgements is implausible, that would further support my general conclusion that not all normative judgements are desires.

cognitivism is supposed to have over cognitivism. But other than in a brief note,² I shall not address such issues: here onwards my focus is simply on conativism: the reader may decide for themselves how my discussion bears on non-cognitivism more broadly.

Despite this restriction of focus, there are three ways in which my objection to conativism will extend to nearby views.

First, conativism is a view about the nature of a mental state: normative judgement. By itself, it says nothing about normative language. Some non-cognitivists defend expressivism, which combines non-cognitivism about normative judgement with a further claim about normative language: expressivists say that the meaning of normative utterances is determined by the state of mind they express. In what follows I continue to focus on the nature of various states of mind rather than the meanings of utterances, but my objection might nonetheless have implications for expressivism insofar as many expressivists incorporate conativism into their view.

Second, quasi-realists are in the non-cognitivist tradition but nonetheless claim that normative judgements are beliefs. Insofar as such views are coherent (see e.g. Dreier 2004), they reconcile these claims by distinguishing two different kinds of belief: full blown beliefs, and states which are beliefs only in some minimalist sense of “belief”. They then claim that normative judgements are beliefs only in the second minimalist sense (see again Dreier 2004, especially 26-9). There is nothing to stop conativists from adopting quasi-realism: they merely need to claim that normative judgements are both desires and beliefs, but the latter only in some minimalist sense. My objection to conativism applies equally to quasi-realist conativism.

Third, some non-cognitivists claim that normative judgements are not desires. But if these authors nonetheless maintain that normative judgements have the same motivational profile as desires, that suffices for my purposes (for this combination of claims, see Blackburn 1998, e.g. 9-10, 13-14, 66, and plausibly Gibbard 1990: 55-6, 75). When I object to conativism below, it is this aspect of it that I focus on, and it is therefore unimportant whether some non-cognitivist

² First, some non-cognitivists might claim that normative judgements are states of *approval* and *disapproval*. But if they also wish to maintain that normative judgements can motivate, it seems as though such non-cognitivists will need to *object* to the Humean view that only desires can motivate us! And once they do that, it is far less clear that there is any real reason to deny that normative judgements are beliefs to begin with. For this reason, it seems more likely that such non-cognitivists are tacitly thinking of states of approval and disapproval as desire-like in the relevant ways, and so are really conativists in disguise.

Second, some non-cognitivists might claim that some normative judgements are desires, but that other such judgements are (non-desire-like) states of approval and disapproval. But if different normative judgements are not even the same state of mind as one another, such non-cognitivists will have an even harder time with another classic problem for non-cognitivism: Making sense of normative disagreement and inconsistency. It is not at all clear how one mental state with one content could be inconsistent with a mental state of a distinct kind with a distinct content.

denies that normative judgements are desires for reasons that are independent of their motivational profile. One way of putting the point is to say that when I object to conativism, I really object to the claim that all normative judgements have the desire-like direction of fit (for this notion, see e.g. Humberstone 1992, Smith 1994: 115). To that extent, my objection applies to any non-cognitivist view that endorses this claim, even if they deny that normative judgements are desires, strictly speaking.³

In short, here I focus on conativism, but my arguments seem likely to extend to many theories in the non-cognitivist tradition, either because they incorporate conativism or else because they incorporate claims that are in the relevant respects close enough to conativism. But I leave open that there might be some better non-cognitivist theories that do not commit to conativism or anything relevantly similar. In that case, my arguments against conativism at least highlight some constraints that any plausible formulation of non-cognitivism must meet.

2. The Motivation Argument

One central argument for conativism is the motivation argument (e.g. Blackburn 1984: 187-9, Blackburn 1998: 70, Gibbard 2003: 11-13, Hare 1952: 1, Stevenson 1937: 16). The basic idea behind the motivation argument is that because normative judgements bear a special connection to motivation, they must be desires. This argument is one of the main weapons that conativists have against rivals, such as cognitivist naturalists. But formulating this argument in a manner that is both precise and plausible has proved difficult. Here I suggest that it is attractive to formulate the argument in roughly the following manner (cf. Snare 1991: 58):

P1) Normative judgements can motivate.

P2) Only desires can motivate.

So, C) Normative judgements are desires.

This argument is still somewhat ambiguous, in that the scope of P1 and C is unclear. In later sections, I shall argue that when we disambiguate the scope of these premises, the argument is either unsound or else fails to support conativism. But before we get to that, I want to first note the virtues of formulating the motivation argument along the above broad lines: the relevant points will not hinge on how we disambiguate the scope of P1 and C. Unlike other formulations of the motivation argument, this formulation focuses on the capacities (powers) of the relevant

³ Do my arguments also extend to “hybrid” non-cognitivist theories, which analyse normative judgements as combinations of desires and beliefs (see e.g. Fletcher and Ridge 2014, Ridge 2014)? So far as I can see, there is no straightforward answer to this question, and it may depend on the hybrid theory in question. But it is worth noting that at least one prominent hybrid non-cognitivist seems to endorse the view that I describe and criticize in section 5 (see Ridge 2014: 19, 177-8).

states of mind.⁴ In this respect, it is parallel to the argument that because H₂O can quench thirst, and only water can quench thirst, water is H₂O. In this kind of argument, note the obvious fact that the premises by no means imply that the relevant powers are always being exercised. This formulation of the motivation argument relies on the claim that normative judgements have the power to motivate, but is consistent with the fact that they often fail to exercise that power, such as when we are weak willed, or make other contrary normative judgements. We will see the significance of this shortly when I contrast this argument with other formulations of the motivation argument.

I take it that P1 of this argument – the claim that normative judgements can motivate – is attractive (cf. Broome 1997: 139, Dancy 1993: 22-3, Ridge 2014: 50). It is attractive to think that people can do things because they think they have good reason to do them, and to repeat, the premise permits that people can be left cold by their normative judgements. Once that is acknowledged, it's hard to see any opposition to the premise that isn't theory-driven. P2 claims that only desires have the capacity to motivate us to act. This is the Humean theory of motivation, in one form. Again, I take it that this claim is attractive. When we explain people's motivations, it seems as though any explanation must ultimately appeal to their desires.

These brief remarks are not intended to demonstrate P1 and P2 above. I like to think that there is at least a good *prima facie* case for accepting them, and in turn a *prima facie* case against views that are inconsistent with the conclusion of the above argument. But adjudicating this dispute is not my main concern in this paper. Rather, I want to address whether these premises can have their scope disambiguated in a manner that allows them to support conativism. So in what follows I shall simply assume that P1 and P2 above are *broadly* plausible, and shall instead focus our attention just on the appropriate scope of P1 and C.

For the purposes of this paper, I treat the Humean theory as the positive claim above, and not as the negative claim that beliefs do not have the capacity to motivate us to act. That negative claim would fit less neatly with quasi-realist conativist views on which normative judgements are beliefs in addition to their being desires. Further, that negative claim is not well supported by the best arguments for the Humean theory. One well-known argument for the Humean theory is Smith's (1994: 116) which appeals to the distinctive direction of fit of desire. If successful, that argument shows that only states with the desire-like direction of fit can motivate, but doesn't directly show that any state with a belief-like direction of fit cannot motivate, since it leaves open that some states have both directions of fit at once.⁵ But the best argument for the Humean theory is simply that such a view is extremely parsimonious and yet has great explanatory power.

⁴ With this in mind, I am treating "can" as a predicate ascribing a power, though a modal version of the argument could also be formulated.

⁵ Of course, Smith does *also* argue that no state can have both directions of fit at once (Smith 1994: 117-125). But that is a further independent claim, and a dubious one at that (see Little 1997: 63-4, Price 1989: 120-1).

It promises to explain all human behaviour by appeal to just two kinds of mental state: desires and means-ends beliefs. That argument obviously appeals to the explanatory power of the claim that only desires can motivate, not to the lack of explanatory power of a theory on which beliefs can motivate, and so supports the positive claim P2 and not the negative claim that I ignore.

Relatedly, we might worry that P2 ought to be formulated as “Only pairs of desires and means-ends beliefs can motivate” (e.g. Smith 1994: 92, Davidson 2001: 3-4). If we formulated it that way, the above argument would be invalid: at best we could conclude that normative judgements are either desires or means-end beliefs. I am not sure whether sympathisers of the Humean theory should be happy with this reformulation of the premise: we might think that desires are the real motivational workers, and that our means-ends beliefs do not themselves partly motivate us, but instead just channel the motivational powers of desire in new directions. But regardless, even if we reformulated this premise in this manner, we could easily reformulate P1 of the argument to maintain the validity of the argument. We could just reformulate P1 to say that normative judgements can motivate, and can do so in a way that means-ends beliefs cannot. After all, the thought driving P1 is that there is some special connection between normative judgement and motivation, and that thought is lost if we permit that normative judgements merely play a role in motivation that ordinary beliefs can also play. For ease, in what follows I shall stick with the simpler formulation of the argument above.

Let me briefly highlight one virtue of my formulation of the motivation argument, by contrasting it with two others. First, a standard way of thinking of the motivation argument has it appeal to the following premise instead of P1 above (e.g. Parfit 2011: 381):

P1*) The judgement that I ought to φ necessarily motivates.

P1* is sometimes called “classic” or “mad dog” judgement internalism (e.g. Björklund et al 2011: 125, Gibbard 2003: 153, Hare 1952: 163-9). But as many have noted, P1* is implausible (e.g. Svavarsdóttir 1999: 176-183). Through weakness of will, we might fail to be at all moved by judgements about what we ought to do. (Here I mean not merely that we might fail to act on some judgement, but that we might fail to have any motivation to act whatsoever, even a motivation that is outweighed.) Since P1* is implausible, it might seem that the motivation argument for conativism cannot be sound. But P1, above, unlike P1*, permits the possibility of weakness of will, and thus evades this objection.

Since P1* is implausible, others have formulated the relevant claim in other more modest ways. One popular alternative says something like the following (e.g. Korsgaard 1986: 15, Smith 1994: 12, 61, van Roojen 2010: 499):

P1**) The presence of the judgement that I ought to φ normally/rationally entails the presence of motivation to φ .

P1** has a definite advantage over P1*, since it permits weakness of will. Perhaps some claim along the lines of P1** is true. But this comes as a hollow victory for fans of the motivation

argument, since it is hard to formulate the motivation argument in a manner that can appeal to this premise and validly get to the conclusion that normative judgements must be desires (Brink 1997: 7-8, Svavarsdóttir 1999: 165-6fn). Since P1** is a claim about mere covariation, it is consistent with the thought that normative judgements influence motivation only through their influence on independent desires, and this is consistent with denying conativism (Smith 1994). P1 is more significant since it makes a claim not about mere covariation, but instead more directly about explanation.

(We could of course modify P1** so that it too is a claim about explanation: we might say, for example, that rational agents are motivated by their normative judgements. Such claims might well be plausible, and could also play a role in the motivation argument. But in that form the relevant claims add nothing to the motivation argument beyond what P1 already provides. As such, the proposed modified version of P1** is unnecessarily bold, requiring us to defend claims (e.g. about rationality) that go beyond the commitments of P1 but which do nothing to make the motivation argument more forceful.)

In short, whilst P1* makes the motivation argument unsound, P1** is likely to make it invalid. In contrast, P1 of the argument as I have formulated it seems modest enough to be plausible but yet bold enough to make the argument valid. To repeat, my goal here was not to show that the argument as I formulate it is conclusive, but rather to show that it has prima facie appeal and can survive standard objections. This is enough to set the scene for the rest of the paper, in which I examine what this argument might show.

3. The Problem: Other-Regarding Normative Judgements

We should ask a simple question about the motivation argument: whether P1 and C are supposed to be read as being universally quantified or not. Since conativism is the view that all normative judgements are desires, conativists should presumably formulate the argument with both claims universally quantified. In that form, let's call it the bold motivation argument. It reads:

Universal-P1) All normative judgements can motivate.

P2) Only desires can motivate.

So, Conativism) All normative judgements are desires.

In this section, I shall first offer some counterexamples to Universal-P1. This suggests that the bold motivation argument fails to establish conativism. The bold motivation argument seeks to establish conativism by appealing to the motivational powers of our normative judgements, but some of our normative judgements do not have any motivational powers. After presenting this argument against the bold motivation argument, I will suggest that the very same counterexamples casts doubt on conativism itself.

It seems as though other-regarding normative judgements have no motivational powers of their own. For example, if Jane judges that Jeff ought to buy his child a birthday present, it seems that this judgement has no power to motivate Jane to do anything. It is a judgement about what Jeff should be doing, and by itself has no bearing at all on what Jane herself will do. When Jane judges that Jeff ought to φ , that is a matter of recognition, not inclination. This casts doubt on Universal-P1.

This point can be obscured by the fact that other-regarding normative judgements can play a role in inference to further normative judgements which do have motivational powers. For example, if Jane *also* judges that she ought to assist others in doing their duty, she might infer that she ought to help Jeff buy the present, and this judgement might motivate her. But here the motivational power is infused only by the addition of a further self-regarding normative judgement: without it, the original other-regarding normative judgement is motivationally inert. So this possibility fails to show that other-regarding normative judgements have motivational powers of their own.

Similar reasoning undermines other possible reasons for endorsing Universal-P1. For example, one might think that if Jane judges that Jeff ought to keep his promises, but finds that he doesn't, this might motivate her to avoid him. Or one might think that if Jane judges that Jeff doesn't invest his money as he ought, this might motivate her to avoid lending him money. But plausibly what really motivates Jane in the first case is the judgement that *she* ought not trust people who don't keep their promises, and what really motivates Jane in the second case is the judgement that *she* ought not lend her money to people who are bad with money. As such it is again doubtful that her other-regarding normative judgement themselves have motivational powers.

The starkest counterexamples to Universal-P1 are those where one person judges that another ought to do something that conflicts with the goals of the first. Most extremely, imagine that Jane is a consistent egoist, who judges that everyone ought only to promote their own wellbeing. Jane might thereby judge that Jeff, in his dealings with her, ought to use and abuse her. This judgement, it seems clear, would have no motivational power over Jane. The point can be seen equally well in less extreme cases: In a prisoner's dilemma, Jane might judge that her partner ought to rat, but it's doubtful that this alone can motivate her to do anything. Or for a final example, I might judge that you ought to save your mother at the expense of mine, but not want you to do so. In general, we will find stark counterexamples to Universal-P1 whenever the relevant norms are believed to be agent-relative, as many prudential and moral norms are believed to be. In such cases, an agent can think that some norm applies to another but not to themselves, and as such be left cold by their recognition of that norm.

One might reply that Universal-P1 says only that normative judgements *can* motivate, and for that reason it is consistent with the fact that they often fail to do so. But the worry is that in cases like those above it is not even plausible that the relevant judgements *could* motivate us:

there are *no* conditions under which other-regarding normative judgements motivate. It just doesn't seem intelligible for someone to act in some way because they recognise that *someone else* ought to do something. Certainly, agents in the cases above are not merely being akratic: It is not as though egoists, or prisoners in the prisoner's dilemma, are being weak-willed when they refuse to help others do what they ought to do. Indeed, if we think that there is a rational requirement not to be akratic, we would thereby commit ourselves to the claim that (e.g.) the prisoner who fails to persuade her partner to rat is being *irrational*, and this is highly implausible. Nothing need be irrational or even abnormal about an agent who judges that *they* have no reason to comply with a norm that they judge governs someone else but not themselves.

Another way of looking at this problem is to remember that the first premise of the motivation argument is a form of judgement internalism. But no plausible formulation of that view makes a claim about all normative judgements. Judgement internalism, as standardly formulated, makes a claim only about self-regarding normative judgements (see references above). So the motivation argument, if it is supposed to support conativism, must appeal to something bolder than judgement internalism: the view that my normative judgements about anyone have motivational power over me. This claim is far from obvious (cf. Gibbard 1990: 100-1, Ridge 2014: 19).

In light of this, it seems reasonable to deny that other-regarding normative judgements have motivational powers, and in turn reasonable to deny Universal-P1 of the bold motivation argument. To that extent the bold motivation argument is unsound and so fails to establish conativism. Perhaps there is some other way of formulating this argument, but conativists would have to show what that formulation is, and show that it avoids commitment to Universal-P1 above. Certainly, we should not assume that there is any straightforward argument from judgement internalism and the Humean theory of motivation to conativism, since judgement internalism is best formulated as a claim about only some specific normative judgements, and need not teach us anything about the nature of the whole class.

So far I have suggested that the motivation argument needs to be reformulated in some non-obvious way if it is to soundly support conativism. Perhaps conativists might simply drop the motivation argument and maintain their view by appeal to other arguments, though we might worry that this amounts to abandoning one of the most persuasive arguments for their view. But there is a further and larger problem: The counterexample to Universal-P1 seems to extend to cast doubt on conativism itself.

If other-regarding normative judgements are motivationally inert, that strongly suggests that they are not desires. Indeed, this straightforwardly follows if we accept the popular theory which analyses desires precisely in terms of their capacity to motivate (see e.g. Smith 1994: 92-129, Stalnaker 1984: 15). It is true that desires are normally thought to motivate only when combined with suitable means-ends beliefs (see pp.5-6, above), but this does nothing help the conativist: it's clear that we do often have the relevant means-ends beliefs and still lack the corresponding motivations. For example, Jane might judge that Jeff ought to use and abuse her, and believe that

she can get him to do so by anonymously sending him the works of Ayn Rand, but still not have any motivation to do so. It seems to follow that her judgement that Jeff ought to use and abuse her is not a desire that he use and abuse her, and it is not clear what other desire it might be.

Perhaps we could add other conditions which are necessary for desires to motivate, or else distinguish between motivating and non-motivating desires (see perhaps Mele 1995). Such claims might allow us to say that other-regarding normative judgements might be desires even if they are motivationally inert. But even if this line of reasoning could be maintained, it is nonetheless independently implausible that there is a necessary connection between other-regarding normative judgements and desire. If Jane judges that Jeff ought to buy his child a birthday present, that seems to leave completely open whether she altruistically hopes he does, spitefully hopes he doesn't, or just doesn't care either way about what she sees as Jeff's business. And again more starkly, a committed egoist surely need not desire that others do what they ought to do, and someone in a prisoner's dilemma can recognise that their opponent ought to rat without desiring that they do so. Even independently of anything we say about motivation, it is not plausible that we necessarily desire that others do what they ought.

In short, when we reflect on other-regarding normative judgements, it seems as though they serve as counterexamples to the claim that all normative judgements have motivational powers, and as counterexamples to the claim that all normative judgements are desires. Reflection on such judgements seems to thereby cast doubt on the motivation argument for conativism, and on conativism itself. When we make judgements about what others ought to be doing, that seems to be a matter of mere recognition, and need not involve any inclination on our own part. In the following two sections, I consider two possible replies open to the conativist. In each case, I argue that the relevant response opens the resulting theory to other objections. We should conclude that the objection above places a significant constraint on any plausible formulation of conativism, and that it is at least unclear whether any independently plausible conativist theory can meet that constraint.

4. Reply 1: Reactive Attitudes

We might think that certain reactive attitudes – such as blame and guilt – are important aspects of morality (e.g. Blackburn 1998: 8-14, Gibbard 1990: 41-5). This might encourage the conativist to claim that other-regarding normative judgements are desire-like after all. In this section I focus on our dispositions to *blame* others for wrongdoing, though I take it that the relevant arguments would extend in obvious ways to nearby alternatives such as our dispositions to shun or to punish wrongdoing, or to praise or reward virtue.

The conativist might press the above line of reasoning above in two different ways. First, they might take the fact that we are disposed to blame others for wrongdoing as evidence that we desire others to act rightly. Second, they might identify our disposition to blame others for

wrongdoing with a desire: they might claim that other-regarding normative judgements just *are* (or involve) desires to blame others under relevant circumstances.

The first of these views is still committed to the counterintuitive claim that people necessarily want others to do what they ought to do. It tries to justify this claim by appeal to our emotional dispositions, but it seems that any such justification will be defeated by the fact it is independently implausible that people *do* necessarily want others to do what they ought to do. That is just what I argued above: it is doubtful that Jane necessarily wants Jeff to buy his child a present, and very implausible that egoists, or people in the prisoner's dilemma, necessarily want others to do what they ought. Even if many of us often do want others to do what they ought, this connection seems highly contingent, and to that extent, we cannot identify other-regarding normative judgements with desires for others to do the relevant things.

But the second of the above views is little better: it is equally implausible that people necessarily want to blame others for wrongdoing. One simple worry is that someone might judge an act to be wrong but blameless: such an agent would endorse an other-regarding normative judgement but have no corresponding desire to blame them (cf. D'Arms and Jacobsen 1994, Ridge 2014: 142-3). But even if we set this aside, it is clear that this second view seems plausible only if we focus solely on moral normativity. It might be true that there is some necessary connection between moral attitudes and dispositions to blame. But to the extent that this is plausible, it seems plausible because we are thinking of morality as a kind of social phenomenon (cf. Gibbard 1990). To that extent, other normative domains which are less social in nature seem to lack any necessary connection with our reactive attitudes. Think, for example, about prudential judgements. Imagine that Jane judges that prudence requires Jeff to buy new running trainers. It is doubtful that she must thereby want him to buy those trainers: she might make this judgement and yet not really care whether he does what he ought. Given that she might not care whether Jeff does what he ought, it seems no huge step to suppose that she might also fail to care what happens to Jeff as a result of his failing to do what he ought (again, cf. Ridge 2014: 143). If she doesn't care much about Jeff at all, she might care neither whether he is prudent nor whether he is chastised when he isn't.

Note that I need not claim that such attitudes are common. The point is simply that it is not a necessary condition on judging that someone else ought to φ that you desire to blame them if they don't. If we find some alien culture where people have no concept of blame, it is far from clear that this would conclusively demonstrate that that same culture has no normative concepts at all: couldn't a utilitarian culture have various views about what people ought to do, but no corresponding inclinations to blame anyone for wrongdoing, considering it counterproductive? Such possibilities seems perfectly coherent, but could not be possible if other-regarding normative judgements just *are* desires to blame others.

In short, it may be plausible that there is *some* connection between some normative judgements – especially moral judgements – and our reactive attitudes, such as our dispositions to blame

others. But it is not plausible that there is a necessary connection between other-regarding normative judgements and our dispositions to blame others, and this casts doubt on this conativist strategy.

5. Reply 2: Gibbard

We might think that other-regarding normative judgements are not directly motivating, but that they nonetheless qualify as desires (or at least desire-like) because they have indirect motivational influence. One initial problem with this suggestion is that many states of mind have some indirect motivational influence. For example, regular beliefs have some influence on what we are motivated to do. So to maintain this suggestion, we need to find some distinctive kind of indirect motivational influence that is had by other-regarding normative judgements but not by other states of mind which the conativist wants to exclude from their theory.

So far as I can see, the best option for the conativist is to claim that other-regarding normative judgements influence motivation in a way that depends only on combining them with prior beliefs. That kind of motivational influence is not had by other states of mind such as regular beliefs. And if we combine this thought with the common thought that any state of mind that can motivate when combined only with belief(s) is by definition a desire (e.g. Smith 1994: 92-129, Stalnaker 1984: 15), we can infer that other-regarding normative judgements are desires.

If Jane's judging that Jeff ought to φ can motivate her by being combined only with a belief about how to get Jeff to φ , then such a judgement just is the desire that Jeff φ . That is the suggestion that we have already rejected: it seems plausible that we can judge that others ought to do things we do not want them to do. The present suggestion is instead better developed in the manner explained by Allan Gibbard (2003: 49-53, see also Gibbard 2012: 174-7, and Ridge 2014: 19, 177-8).⁶ On Gibbard's view, other-regarding normative judgements are desires about what to do if you were in the other person's place. That is, Gibbard claims that Jane's judgement that Jeff ought to φ is a conditional desire to φ if she were in Jeff's place. Such a view rightly permits that Jane might judge that Jeff ought to φ but not want him to, and yet nonetheless treats such judgements as desires, as conativists must. And such conditional desires do indeed have only indirect motivational potential, as promised: conditional desires need to be combined with the belief that the relevant condition is met if they are to motivate. As applied in our case, we get the conclusion that Jane's judgement can motivate her, but only by being combined with the belief that she is Jeff.

We should distinguish two more precise ways we might develop Gibbard's proposal. First, there is the simpler option:

⁶ Gibbard treats normative judgements as *plans* rather than desires, but as I said in section 1, this small kind of difference seems unimportant for our purposes, since Gibbard nonetheless treats plans as being like desires in the way that they motivate. For simplicity, I shall continue to talk in terms of desire.

Gibbard-simple: When A judges that B ought to φ , that consists in A's conditionally desiring to φ if they, A, were in B's circumstances.⁷

Second, there is the more sophisticated option:

Gibbard-sophisticated: When A judges that B ought to φ , that consists in A's conditionally desiring to φ if they were B and in B's circumstances.

The difference between the proposals is that the sophisticated proposal, unlike the simple, understands other-regarding normative judgements to consist in desires for circumstances where we have different haecceities than we in fact have. So far as I can tell, Gibbard himself endorses the second option (2003: 50-1), but for completeness I shall object to both, in turn.

So first, there is Gibbard-simple. The first thing to note is that Gibbard-simple is counterintuitive. It is counterintuitive to say that when Jane makes a normative judgement about Jeff, she is forming a conditional desire for the eventuality that she end up in his circumstances. The view is all the more surprising once we remember that "circumstances" here has to include not only Jeff's external environment, but also anything that might be relevant to what he ought to do: his ignorance, his character traits, his emotions, and so on (cf. Gibbard 2003: 50-1, 2012: 174-5). So the suggestion had better be that when Jane judges that Jeff ought to φ , this is a desire of hers to φ if she had all of his properties. And that is a conditional desire for a possibility that is extremely unlikely to occur, if it's possible at all.⁸ Other-regarding normative judgements are a familiar everyday mental state, and it would be surprising if they turned out to be an attitude directed towards such bizarre counterfactual circumstances; other-regarding normative judgements certainly do not *seem* to be attitudes towards extremely remote possible worlds, but instead attitudes directed towards other people in the actual world.

One way to think about this is to think about how children learn to make normative judgements about others: it is highly doubtful that they do so by thinking about various highly remote possibilities. If a young boy thinks that girls aren't supposed to have short hair, they do not come to that thought by imagining themselves as a girl (indeed, a failure to do so is presumably part of the problem). *Perhaps* it might be conceded that children *sometimes* form views about others by considering the relevant counterfactuals, but it is highly doubtful that they *always* do so, especially

⁷ This sort of view may also have been held by Hare (see e.g. 1981, especially chapter 7). But matters are not so clear because Hare often talks about what a person is *committed to*, and this more often sounds like a normative claim rather than a descriptive one. Gibbard-simple says that other-regarding normative judgements literally *are* desires of the relevant sort, and Hare may have meant to commit only to the weaker thesis that other-regarding normative judgements *rationaly require* desires of the relevant sort. I say nothing here against this latter thesis, which might well be true, but fails to provide the conativist with what they need. See also the discussion of supervenience, below.

⁸ In passing, Gibbard claims that even if it's sometimes *metaphysically* impossible for one person to be in another's exact circumstances, it's nonetheless *epistemically* possible (2012: 177). But even if this is right, it's unclear how this is supposed to dispel the oddness of the view given how epistemically remote those possibilities often are.

when the relevant counterfactual possibilities are (or are seen to be) very remote. So the first concern is that Gibbard-simple is highly counterintuitive, and to that extent it seems *ad hoc*.

Gibbard-simple also faces a second objection. Gibbard-simple analyses Jane's judgement [that Jeff ought to φ] as Jane's desire [to φ if she were in Jeff's circumstances].⁹ Presumably, Gibbard-simple would also have us analyse Jane's judgement [that she ought to φ if she were in Jeff's circumstances] as Jane's desire [to φ if she were in Jeff's circumstances]. Since these two judgements receive the same analysis, Gibbard-simple forces us to identify them. That is, according to Gibbard-simple, there is no difference between Jane's judgement [that Jeff ought to φ] and Jane's judgement [that she ought to φ if she were in Jeff's circumstances]. But the problem is that these two judgements are distinct. This is clearest when we consider the obvious fact that Jane might have concluded that Jeff ought to φ without having extended the conclusion to her own case. The reverse is also possible: Jane might have made a plan for herself in Jeff's circumstances without having actually considered what Jeff himself ought to do. By analysing other-regarding normative judgements as conditional desires, we lose the ability to give an independent analysis of genuinely counterfactual normative judgements.

Might Gibbard reply that other-regarding normative judgements and genuinely counterfactual normative judgements are in fact the same state of mind, and claim that this truth is merely non-obvious to us? But that reply suggests that there is no such thing as an inference between these states of mind, and that is implausible. Plausibly, we can sometimes get people to change their minds about how they judge others precisely by having them consider the relevant counterfactual claims about themselves, and vice versa, and this would not be possible if the relevant states of mind were literally identical.

Gibbard-simple gains illusory plausibility here because normative truths supervene on non-normative truths (henceforth: Supervenience). With Supervenience in mind, it is tempting to think that if Jeff ought to φ , then Jane ought to φ if she were in Jeff's exact circumstances. But even if Supervenience is true, that doesn't tell us much about Jane's judgements. Jane might fail to accept Supervenience, or more likely, might fail to accept every single implication of that truth. So Supervenience does not show that Jane's judging that Jeff ought to φ is the very same thing as Jane's judging that she ought to φ if she were in Jeff's circumstances.

Some non-cognitivists have claimed that Supervenience is not a metaphysical truth, but instead a conceptual one (e.g. Blackburn 1993b). But even this will not help. Even if Jane's judgements embody failures to accept conceptual truths, this is no bar to her having those judgements (cf. Williamson 2007: 73-133).¹⁰ Jane might be conceptually confused and deny Supervenience. Or

⁹ Here, and later, I sometimes use square brackets to mark the contents of attitudes.

¹⁰ At one point Blackburn seems to assert the reverse (1993a: 122) but he gives no argument for this bold claim, and it isn't required for the supervenience argument that he presents against moral realism.

again, more likely, she might accept Supervenience but fail to accept various truths that are entailed by combining Supervenience with other beliefs that she holds – she might fail to combine her very abstract commitment to Supervenience with her judgement that Jeff ought to φ (and it is surely possible to fail to endorse every implication of a conceptual truth one recognises).

In short, Gibbard-simple is counterintuitive, and moreover it collapses the distinction between Jane's judging that Jeff ought to φ and Jane's judging that she ought to φ if she were in Jeff's position. Those judgements are distinct, and Supervenience does not show otherwise.

I now turn to Gibbard's own preferred view, Gibbard-sophisticated. To remind you, it says:

Gibbard-sophisticated: When A judges that B ought to φ , that consists in A's desiring to φ if they were B and in B's circumstances.

It's not obvious that Gibbard-sophisticated really improves on Gibbard-simple. It does allow us to distinguish Jane's judgement [that Jeff ought to φ] and Jane's judgement [that she ought to φ if she were in Jeff's circumstances]. Gibbard-sophisticated entails that the former but not the latter consists in a desire that is conditional on Jane's being Jeff. But this is not obviously progress, since we might now worry that Gibbard-sophisticated fails to distinguish a different pair of judgements: Jane's judgement [that Jeff ought to φ] and Jane's judgement [that she ought to φ if she were Jeff and in Jeff's circumstances]. Gibbard-sophisticated will presumably give these two judgements the same analysis, and this might seem mistaken for just the same reasons as those given above. Again, by analysing other-regarding normative judgements as conditional desires, we lose the ability to give an independent analysis of genuinely counterfactual normative judgements.

That said, here it is admittedly *somewhat* less clear-cut that these two judgements really are distinct. But whatever plausibility Gibbard-sophisticated gains here, it loses with respect to the first worry above. To whatever extent Gibbard-simple was counterintuitive, Gibbard-sophisticated is worse. Gibbard-sophisticated says that people who make judgements about what other people ought to be doing are forming desires for circumstances where their identity differs. It is highly counterintuitive to suppose that we have conditional desires for such impossible circumstances: even if such conditional desires are possible, it does not seem that they are commonly occurring parts of our mental lives. And again, it is deeply implausible that we learn to make other-regarding normative judgements by learning to think about how our own identity and circumstances might differ: it is not clear that the average child is even capable of thinking about such matters. These implications of Gibbard-sophisticated should make us very wary of accepting it unless there is no other option.

I conclude that Gibbard's strategy is at best ad hoc and counterintuitive, and at worst inconsistent with clear distinctions between different normative judgements. I also argued above that conativists should not try to rescue their view by appeal to the connection between

normative judgements and reactive attitudes. With no other obvious option on the table, I conclude that the objection stands and conativism is implausible.

6. The Frege-Geach Objection

It might be worth very briefly comparing the objection I have raised against conativism with the Frege-Geach objection to non-cognitivism (see Schroeder 2008 for comprehensive discussion). As applied to conativism, the Frege-Geach objection is that normative judgements with logically complex contents cannot be analysed as desires.¹¹ We can illustrate this claim by appeal to normative judgements with negated contents (Schroeder 2008: 39-55, Unwin 1999, 2001). Jane's judgement that it's not the case that she ought to drink tea seems hard to analyse as any desire. It is not a desire to drink tea, nor a desire not to drink tea, nor a failure to desire to drink tea. Worse, even if we could find some way to analyse this judgement as one of these desires, that would only move the problem elsewhere: the judgement more naturally associated with the relevant desire would itself now lack a suitable analysis. That is, there is no way to reduce the following four states of mind on the left in terms of the three on the right:

Judgement that she ought to φ ($JO\varphi$)	Desire to φ ($D\varphi$)
Judgement that she ought not φ ($JO\neg\varphi$)	Desire not to φ ($D\neg\varphi$)
Judgement that it's not the case that she ought to φ ($J\neg O\varphi$)	???
Failure to judge that she ought to φ . ($\neg JO\varphi$)	Failure to desire to φ ($\neg D\varphi$)

How does this objection relate to the objection I have raised in this paper? The two are distinct, insofar as the examples I have focused on have been logically atomic normative judgements (e.g. the judgement that [Jeff ought to φ]), and the Frege-Geach objection appeals to logically complex normative judgements. But though these two problems are distinct, they point to the

¹¹ Often, the Frege-Geach objection is expressed as the problem for expressivism of accounting for the meaning of logically complex sentences that employ normative predicates. But since I have defined conativism as a theory about states of mind, rather than the meanings of sentences, the Frege-Geach objection applies to conativism only if we express it – as I do in the main text – as an objection that makes reference to the nature of states of mind rather than to the meanings of sentences (this is not wholly unusual, see e.g. Unwin 1999, 2001). This is plausibly the best way to think about the fundamental source of the Frege-Geach problem: Expressivists claim that meanings are inherited from the states of mind they express, and so if we can find states of mind that constitute logically complex normative judgements, it seems likely that we thereby give expressivists the materials they need to explain the meanings of logically complex normative sentences.

Interestingly, Mark Kalderon takes the reverse view (2005). He too distinguishes between the psychological theory of non-cognitivism, and the semantic theory of expressivism (2005:52-3, and 95-146). But in contrast to me, he claims that the Frege-Geach problem is generated by expressivism rather than non-cognitivism (2005: 52-94). It is for this reason that he claims that his *Fictionalist* theory, which commits to non-cognitivism but not expressivism, avoids the Frege-Geach problem. But the claim that Kalderon's brand of Fictionalism avoids the Frege-Geach problem is mistaken – see Eklund 2009 and references therein.

same conclusion: that not all normative judgements are desires (cf. Finlay 2014: 130-134). Perhaps there are other respects in which conativism is problematic, but the objection I have raised, and the Frege-Geach objection, both focus on one specific feature of conativism: that it analyses all normative judgements as desires. Any theory that does not commit to this claim promises to thereby avoid both objections.

7. Other Theories

I have suggested that there are some grounds for doubt about the motivation argument for conativism, and equal grounds for doubt about conativism itself: other-regarding normative judgements do not seem to have motivational powers, and do not seem to be desires. But this leaves open that we might accept some other more modest formulation of the motivation argument and some other more modest conclusion.

The problematic judgements that I focused on were other-regarding normative judgements. So we might formulate the motivation argument with reference to self-regarding (de se) normative judgements and draw the conclusion that such normative judgements are desires. But we might restrict the motivation argument further. For example, given the Frege-Geach problem, we might doubt that logically complex normative judgements have motivational powers, and doubt that they are desires. So we might formulate the motivation argument with reference to self-regarding logically atomic normative judgements, and draw the conclusion that just those judgements are desires. In fact, I am going to restrict the argument still further, and formulate it with reference to normative judgements with the content [I have reason to φ]. This formulation of the motivation argument is more restricted than my arguments warrant. I work with this formulation of the motivation argument primarily for simplicity: the claims that follow are easier to understand if we express them in these simple and positive terms. As it happens, I also think that this is the most plausible way of specifying the class of normative judgements with motivational import, but I shall not rely on or defend that claim here.

That is, in what follows, I address the following argument, which I label *the best motivation argument*:

P1+) All normative judgements with the content [I have reason to φ] can motivate.

P2) Only desires can motivate.

So, C+) All normative judgements with the content [I have reason to φ] are desires.

Obviously, P1+ is consistent with the claim that other-regarding normative judgements cannot motivate. So too, C+ is consistent with the claim that other-regarding normative judgements are not desires. So the best motivation argument avoids the problems facing the bold motivation argument. But at the same time, the best motivation argument fails to support conativism, which makes a claim about all normative judgements, not just some normative judgements with certain specific contents.

If we think the best motivation argument looks compelling, we might be inclined to reject pure cognitivism, which claims that normative judgements are never desires. Such a view requires that we reject P1+ or P2 above, and those claims seem attractive. But I have also suggested that we should reject conativism, which claims that all normative judgements are desires. So the remaining possibility is that some normative judgements are desires, and some are not. Such a view would be supported by the best motivation argument, and would not be threatened by other-regarding normative judgements.

There may be very many views of this kind (see, possibly, Little 1997, McDowell 1998, McNaughton 1988: 106-117, Pettit 1987, Price 1989, and Scanlon 1998: 7-8, 37-49). Here I describe just two, to illustrate the kind of view I have in mind. First, we might adopt some kind of desire-as-belief view, such as the following (see REDACTED, cf. Humberstone 1987, McNaughton 1988: 106-117):

Desire-as-belief: To desire to φ just is to believe that you have normative reason to φ .

Desire-as-belief reduces desires to a particular kind of normative belief. According to desire-as-belief, normative judgements about oneself having reason to do things are desires, but other normative judgements are not desires. Whatever else we might say about desire-as-belief, it should be clear that it would have the attractive features we want. Desire-as-belief is supported by the best motivation argument, since it explains how the relevant normative judgements can motivate us: because they are desires. It also fits well with the examples I have discussed in this paper, since it says that those normative judgments cannot motivate us and are not desires. Desire-as-belief seems superior to pure cognitivism because it is consistent with the premises of the best motivation argument, and superior to conativism because it is consistent with the fact that other-regarding normative judgements are not desires.

(Given the focus of the paper, I here focus on the advantages of desire-as-belief with respect to other-regarding normative judgements. But given that I above said that this objection to conativism is structurally similar to the Frege-Geach objection, it should be clear that desire-as-belief promises to avoid both problems. Desire-as-belief says that only (some) logically atomic normative judgements are desires, and thereby rightly entails that no logically complex normative judgement can motivate or is a desire. It thereby avoids the Frege-Geach objection. If the best argument for conativism is the motivation argument, and the worst objection the Frege-Geach objection, Desire-as-belief is clearly the superior view: It has the former advantage but not the latter cost.)

It is worth very briefly reminding ourselves of one of the main arguments for desire-as-belief (see e.g. Quinn 1995, REDACTED). This argument says that we have to adopt desire-as-belief if we want to explain how it is that desires rationalise our behaviour. If we think of desires as mere drives, pushing us around, they seem to explain our physical movements in a mechanistic way that does not show our behaviour to be rational. In contrast, if we think that our desires

represent their objects as normatively favoured in some way, it seems that we can explain why it is that desires rationalise, rather than merely cause, our behaviour.

This argument not only promises to justify desire-as-belief, but also helps us to respond to a possible objection to the view that arises in the context of this paper. One might worry that desire-as-belief merely *stipulates*, rather than *explains*, the fact that other-regarding normative judgements are not desires. But if desire-as-belief is justified by appeal to the argument above, we can reply to this objection. The argument above suggests that desires are privileged in explaining our behaviour not because only they can cause physical movements, but because only they can rationalise, rather than merely cause, our behaviour. But if this is so, we can immediately see why other-regarding normative judgements are not desires: because their normative content could not rationalise our own behaviour. To this extent, desire-as-belief does promise to explain why other-regarding normative judgements are not desires. No doubt more could be said here, but it is clear that if desire-as-belief is attractive at all, it has the resources to explain why other-regarding normative judgements are not desires.

It's helpful to approach the overall topic from another direction. Conativism and desire-as-belief might look like very similar theories: Conativism reduces normative judgements to desires, and desire-as-belief reduces desires to normative judgements. So the two views might seem similar in that both identify normative judgements and desires with one another. One issue that divides the two views is whether both normative judgements and desires are beliefs, as desire-as-belief says, or whether neither normative judgements nor desires are beliefs, as conativism says. This issue may be difficult to resolve, especially once conativists adopt quasi-realism and claim that normative judgements are beliefs, in some minimalist sense (Dreier 2004). But there is a further and clearer contrast between the two views, and that is the scope of the identity that they posit. Conativism identifies all normative judgements with desires, whereas desire-as-belief identifies only some normative judgements with desires. Here, I suggest that desire-as-belief has the upper hand. For example, one central difference between desire-as-belief and conativism is that desire-as-belief does not treat other-regarding normative judgements as desires. It can thereby better accommodate other-regarding normative judgements, whilst retaining the central attraction of explaining the motivational power of (some) normative judgements.¹²

A second view that says that some but not all normative judgements are desires is what I'll call *modest* desire-as-belief. Such a view distinguishes between two different kinds of desire (cf. Nagel 1970, Schueler 1995), and analyses only some as normative beliefs:

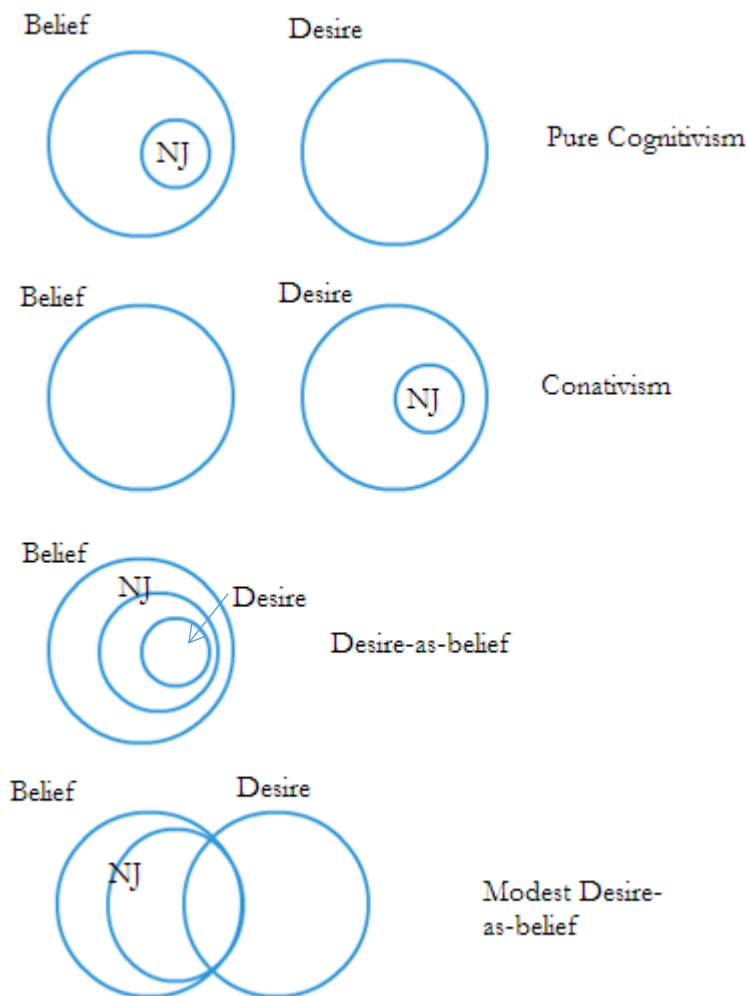
¹² Note also that desire-as-belief permits that even some normative judgements with the content [I have reason to φ] do not motivate us, and thereby permits weakness of will. It commits only to the claim that such judgements have the *power* to motivate us (see REDACTED).

Modest desire-as-belief: Desires are either reflective, or brute. To reflectively desire to φ is to believe that you have a normative reason to φ . To brutally desire to φ involves no normative beliefs.

Modest desire-as-belief says that we have two kinds of desire, and that some reduce to beliefs about what we have reason to do, and others do not. Such a view also gains support from the best motivation argument. It explains how some normative judgement can motivate us: because they are (reflective) desires. But it also permits that other normative judgements cannot motivate us, as I have argued. In these respects, it is just the same as desire-as-belief. It differs from desire-as-belief in that it permits that not all of our desires involve normative judgements, and this might seem attractive when we think about desires for the bad, appetites, and the desires of animals (see e.g. Stocker 1979, Velleman 1992, for responses, see e.g. Raz 1999, 2010, REDACTED).

Perhaps there are other views that also share these merits: that entail that some but not all normative judgements are motivational. Like conativism, such views permit that (some) normative judgements have motivational powers that only desires have. But unlike conativism, they also permit that other normative judgements lack those motivational powers. That is, such views, like conativism, have the payoff that they explain the motivational role of normative judgement. But unlike conativism, they avoid committing to the claim that all normative judgements have such motivational powers.

It might be helpful to illustrate the four candidate views as Euler diagrams (using “NJ” as shorthand for “normative judgement”):



The first Euler diagram represents pure cognitivism. On that view, there is an exclusive distinction between beliefs and desires, and normative judgements are a subset of our beliefs. The best motivation argument threatens to undermine this view, since it threatens to show that some normative judgements are desires. The second Euler diagram represents conativism. On that view, there is an exclusive distinction between beliefs and desires, and all normative judgements are a subset of our desires.¹³ In this paper, I have suggested that this view is also mistaken, since it wrongly treats all normative judgements as desires. That claim is not demonstrated by the best motivation argument, and faces objections from other-regarding normative judgements (as well as from logically complex normative judgements).

The third Euler diagram represents desire-as-belief. According to desire-as-belief there is no exclusive distinction between beliefs and desires. Rather, our desires are a subset of our beliefs.

¹³ Quasi-realist conativists might say that the distinction between beliefs and desires is not exclusive, because normative judgements are both – they are beliefs in some minimalist sense. But this makes no real difference here: such a view is still committed to the claim that all normative judgements are desires.

In particular, they are beliefs with a particular normative content. Since these beliefs are desires, desire-as-belief fits well with the best motivation argument: these beliefs can motivate even though only desires can motivate. But not just any belief with a normative content is a desire: that is how the view is consistent with the examples that limit the scope of the motivation argument, and that cast doubt on conativism. The fourth Euler diagram represents modest desire-as-belief. According to modest desire-as-belief, there is no exclusive distinction between beliefs and desires. Rather, some normative judgements are both. That is how modest desire-as-belief fits well with the best motivation argument: it explains how some normative judgements can motivate us. But again, modest desire-as-belief does not say that all normative judgements are desires, and as such avoids the objections I have been pressing against conativism.

8. Conclusion

In this paper I first suggested that the motivation argument fails to support conativism. Only some normative judgements have the capacity to motivate, and this provides inadequate support for the claim that all normative judgements are desires. Second, I suggested that those same problematic normative judgements cast doubt on conativism itself, since such judgements are not plausibly analysed as desires. Third, I suggested that other views, such as desire-as-belief, are well-placed to accommodate the motivational import of normative judgement. They are consistent with the attractive claims that some normative judgements have motivational powers and are desires. But unlike conativism, they allow that not all normative judgements have motivational powers or are desires. I conclude that such theories may hold greater promise than conativism.¹⁴

Alex Gregory
The University of Southampton
Department of Philosophy
a.m.gregory@soton.ac.uk

¹⁴ For useful advice on this paper, I should thank *at least*: colleagues at Southampton (especially Jonathan Way), audiences at Antwerp, Humboldt and Leeds Universities, Teemu Toppinen, and two referees for this journal.

Bibliography

- Blackburn, S. 1984. *Spreading the Word*. Oxford University Press
- Blackburn, S. 1993. *Essays in Quasi-Realism*. Oxford University Press
- Blackburn, S. 1993a. "Moral realism" in his 1983 pp.111-129
- Blackburn, S. 1993b. "Supervenience Revisited" in his 1983 pp.130-148
- Blackburn, S. 1998. *Ruling Passions*. Oxford University Press
- Brink, D. 1997. "Moral Motivation" in *Ethics* 108 pp.4-32
- Björklund, F., Björnsson, G., Eriksson, J., Olinder, R., and Strandberg, C. 2012. "Recent Work On Motivational Internalism" in *Analysis* 72:1 pp.124-137
- Broome, J. 1997. "Reasons and Motivation" in *Proceedings of the Aristotelian Society Supplementary Volume* 71 pp.131-46
- D'Arms, J. and Jacobson, D. 1994. "Expressivism, Morality, and the Emotions" in *Ethics* 104:4 pp.739-763
- Dancy, J. 1993. *Moral Reasons*. Blackwell
- Davidson, D. 2001. "Actions, Reasons, and Causes" in his *Essays on Actions and Events*. Oxford University Press pp.3-19
- Dreier, J. 2004. "Meta-ethics and the problem of creeping minimalism" in *Philosophical Perspectives* 18 pp.23-44
- Eklund, M., 2009. "The Frege-Geach Problem and Kalderon's Moral Fictionalism" in *Philosophical Quarterly* 59 pp.705–12.
- Finlay, S. 2014. *A Confusion of Tongues*. Oxford University Press
- Fletcher, G. and Ridge, M. (eds) 2014. *Having It Both Ways: Hybrid Theories in Meta-Normative Theory*. Oxford University Press
- Gibbard, A. 1990. *Wise Choices, Apt Feelings*. Harvard University Press
- Gibbard, A. 2003. *Thinking how to live*. Harvard University Press
- Gibbard, A. 2012. *Meaning and Normativity*. Oxford University Press
- Hare, R. 1952. *The Language of Morals*. Oxford University Press
- Hare, R. 1981. *Moral Thinking*. Oxford University Press
- Humberstone, L. 1987. "Wanting as Believing" in *Canadian Journal of Philosophy* 17 pp.49-62
- Humberstone, L. 1992. "Directions of Fit" in *Mind* 101: 401 pp.59-83
- Kalderon, M. 2005. *Moral Fictionalism*. Oxford University Press
- Korsgaard, C. 1986. "Scepticism about practical reason" in *The Journal of Philosophy* 83:1 pp.5-25
- Little, M. 1997. "Virtue as Knowledge: Objections from the Philosophy of Mind" in *Nous* 31:1 pp.59-79
- McDowell, J. 1998. "Are Moral Requirements Hypothetical Imperatives?" in his *Mind, Value, and Reality*. Harvard University Press. pp.77-94
- McNaughton, D. 1988. *Moral Vision: An Introduction to Ethics*. Blackwell
- Mele, A. 1995. "Essentially Motivation Constituting Attitudes" in *The Philosophical Review* 104:3 pp.387-423

- Nagel, T. *The Possibility of Altruism*. Princeton University Press (1970)
- Parfit, D. 2011. *On What Matters, Volume Two*. Oxford University Press.
- Pettit, P. 1987. "Anti-Humeans and Motivation" in *Mind* 96:384 pp.530-5
- Price, H. 1989. "Defending Desire As Belief" in *Mind* 98:389 pp.119-127
- Quinn, L. 1995. "Putting Rationality in its Place" in Hursthouse, R., Lawrence, G. and Quinn, W. (eds) *Virtues and Reasons: Philippa Foot and Moral Theory*. Oxford University Press pp.181-208
- Raz, J. 1999. "Agency, Reason, and the Good" in his *Engaging Reason; On the Theory of Value and Action*. Oxford University Press. pp.22-45
- Raz, J. 2010. "On The Guise of The Good" in Tenenbaum, S. (Ed.) *Desire, Practical Reason, and the Good*. Oxford University Press. pp.111-137
- Ridge, M. 2014. *Impassioned Belief*. Oxford University Press
- Scanlon, T. 1998. *What We Owe to Each Other*. Harvard University Press
- Schroeder, M. 2008. *Being For*. Oxford University Press
- Schueler, G. 1995. *Desire*. MIT Press
- Smith, M. 1994. *The Moral Problem*. Blackwell
- Snare, F. 1991. *Morals, Motivation and Convention*. Cambridge University Press
- Stalnaker, R. 1984. *Inquiry*. MIT Press
- Stevenson, C. 1937. "The Emotive Meaning of Ethical Terms" in *Mind* 46:181 pp.14-31
- Stocker, M. 1979. "Desiring the Bad: An Essay in Moral Psychology" in *The Journal of Philosophy* 76:12 pp.738-753
- Svavarsdóttir, S. 1999. "Moral Cognitivism and Motivation" in *The Philosophical Review* 108:2 pp. 161-219
- Unwin, N. 1999. "Quasi-realism, negation and the Frege-Geach problem" in *The Philosophical Quarterly* 49:196 pp.337-352
- Unwin, N. 2001. "Norms and negation: A problem for Gibbard's logic" in *The Philosophical Quarterly* 51:202 pp.60-75
- Van Roojen, M. 2010. "Moral rationalism and rational amorality" in *Ethics* 120 pp.495-525
- Velleman, D. 1992. "The Guise of the Good" in *Nous* 26 pp.3-26
- Williamson, T. 2007. *The Philosophy of Philosophy*. Oxford University Press