

This is a post-peer-review, pre-copyedit version of an article forthcoming in *Synthese*. Please cite the published version when available.

WEIGHING AIMS IN DOXASTIC DELIBERATION

Short Abstract: In this paper, I defend teleological theories of belief against the exclusivity objection. I argue that multiple epistemic aims interact in concrete cases of doxastic deliberation. I also survey some of the recent literature concerned with this debate.

INTRODUCTION

Many philosophers are tempted by the idea that beliefs aim at the truth. The hope is that this aim hypothesis can provide a unified account of both the descriptive and normative features of belief. At the descriptive level, the suggestion is that beliefs are essentially characterised by their aim: among the propositional attitudes, beliefs uniquely aim at the truth. While at the normative level, the aim accounts for the rationality of belief: if beliefs aim at the truth, then true beliefs are correct because they achieve belief's aim, and believing for epistemic reasons is rational because epistemic reasons indicate truth. However, a question arises about what exactly it means for beliefs to aim at the truth. One influential suggestion, which is the focus of this paper, is that the *aim* of belief is just what it claims to be: a genuine aim (or *telos*) that is essential to belief. This thesis—call it the teleological thesis—holds that believing is, in some important sense, a purposive or aim-motivated behaviour. Just like any other aim-

motivated behaviour, forming and abandoning beliefs is an instrumental process with an aim that can (or can fail to) be satisfied.¹

Now, there are a number of reasons to support this teleological reading of belief's aim, but a couple of important points are worth mentioning. First, in contrast to the theory's main competitor, the teleological thesis promises to provide a wholly naturalistic account of belief. The alternative is to interpret the "aim" of belief as a normative requirement rather than as a genuine aim.² The worry with this alternative, however, is that we must appeal to a somewhat mysterious realm of normativity. This might ultimately be plausible, but the teleologists hope to do the same work without such an appeal. Belief's aim is instead reduced to the mechanisms involved in belief formation and maintenance, such as the intentional processes that come into play when we deliberate what to believe.

Second, the teleological thesis promises to satisfy both the descriptive and normative ambitions of belief theorists. The mechanisms involved in beliefs aiming at the truth serve to demarcate beliefs from other propositions attitudes; and the aim sets belief's standard of correctness. When we form beliefs, there is an end to be achieved (i.e., to believe the truth); and there are appropriate ways (or means) of achieving that end (i.e., to form and maintain beliefs for epistemic reasons).

¹ McHugh (2011, 2012), Sosa (2010), Steglich-Petersen (2006, 2009), and Velleman (2000) all defend teleological theories of belief. Velleman now prefers a normative reading of belief's aim (see Shah and Velleman 2005).

² I am not here primarily concerned with normative theories of belief. I mention them only as a point of contrast. For two influential defences of the normative approach, which focus specifically on the contrast between teleological and normative readings of belief's aim, see Shah (2003) and Shah and Velleman (2005). For a general collection of readings on the topic see Chan (2013).

Yet, despite this initial appeal, the teleological thesis must confront a serious objection—one that, following Steglich-Petersen (2009), I call the *exclusivity objection*. The problem is that reasoning about what to believe does not proceed as we should expect a genuinely aim-motivated behaviour to proceed: the supposed “aim” of belief fails to resemble a genuine aim. As David Owens (2003) explains, properly speaking, an essential feature of aims is that they can *interact* (or be *weighed*) with other aims in deliberation. Typically, in deliberation, we can take multiple aims into account, and choose a course of action that is best all things considered—one that is intended to maximise expected utility. For instance, suppose that you aim to watch a movie tonight, but that you also aim to go to the gym (that same evening). Given that you have both of these aims, you can weigh them in deliberation, and you can choose a course of action that you think is best all things considered. You can decide, say, to spend less time at the gym, so that you can get home in time to watch the movie. In these situations, deliberation involves multiple aims, such that we can decide on (what we take to be) the most appropriate course of action given those aims. However—and this is the problem—deliberating about what to believe (doxastic deliberation) does not appear to permit interaction between belief’s aim and any other aims that we might have. If belief has an aim, then it should be weighable with other aims, but it is not, so the putative aim of belief is not an aim after all.³ In practice, we bear witness to this problem when we observe the difficulty of forming and abandoning beliefs for reasons that do not bear on the truth of a proposition (i.e. for non-epistemic reasons). For instance, you cannot decide to believe that you have just won the lottery, if you haven’t, no

³ It’s important to bear in mind that this debate is framed in the context of doxastic deliberation, which is a *conscious* process. No one denies that non-epistemic reasons can subconsciously influence what we believe. The obvious example is *wishful thinking* (Shah 2003).

matter how much you might desire to hold that belief. We cannot, in this sense, just decide to believe a falsehood, regardless of what our other aims might be. In this way, belief's aim is seemingly unable to be weighed with any broader concerns that we might have, so it's difficult to see how it can be a genuine aim.⁴

At a more fundamental level of analysis, this objection derives from a widely accepted thesis about the nature of belief, known as the *exclusivity thesis*. Many agree that the processes involved in doxastic deliberation are exclusively sensitive to epistemic reasons. The thought is that only reasons which bear on the truth of a proposition can, in principle, influence what we come to believe in deliberation.⁵ If this is true, then weighing belief's aim with other aims begins to look like a conceptual or metaphysical impossibility. There just cannot be interaction between belief's aim and other aims, since forming and abandoning beliefs is not a process that is sensitive to our broader aims. Thus, to summarise the objection:

P1. If an activity is aim-motivated, then the aim involved must be able to interact with other aims in deliberation.

P2. Doxastic deliberation is exclusively sensitive to epistemic reasons (the exclusivity thesis).

P3. Doxastic deliberation does not permit interaction between multiple aims.
[From P2]

⁴ A similar objection is also raised by Kelly (2003).

⁵ Appeals to the exclusivity thesis are apparent on both sides of the debate. For instance, see Archer (2015), Owens (2003), Shah (2003), Steglich-Petersen (2006, 2009), Sullivan-Bissett & Noordhof (2013), Sylvan (2016), and Velleman (2000).

C. Doxastic deliberation is not an aim-motivated activity (the teleological thesis is false). [MT, P1, P3]

In this paper, my primary goal is to defend the teleological thesis against this objection. I do so by arguing that, in specific instances (or concrete cases) of doxastic deliberation, multiple distinct *epistemic* aims interact. This is just to say that P2 does not entail P3, so the objection fails. I begin in Section 1 by outlining a general schema for weighing aims. I then use this schema as a framework for my discussion—it is particularly helpful for delineating the logical space of my position in relation to other recent defences of the teleological thesis, and a brief survey of these other positions is also presented. In section 2, I then develop and defend my positive proposal.

1. WEIGHING AIMS IN DELIBERATION: A SCHEMA

Various kinds of aims can be weighed in deliberation. One broad distinction we can draw is between *epistemic aims* and *pragmatic aims*. In the first instance, epistemic aims are those that have truth as the content of their goal. These aims can motivate and guide actions, such as inquiry and reasoning, and they can be essential features of propositional attitudes, such as guessing (Owens 2003), and, perhaps, belief. In contrast, pragmatic aims are those that do not have truth as the content of their goal. For example, aiming to watch a movie tonight and aiming to go to the gym are pragmatic aims. Instead of seeking to represent the way things are, pragmatic aims involve bringing about certain states of affair. In this way, the distinction between epistemic and pragmatic aims roughly parallels the *mind-to-world* and *world-to-mind* directions of fit, respectively. Epistemic aims involve attempting to represent the world as it is, while pragmatic aims involve attempting to make the world a certain way.

A second distinction we can draw is between *the ways in which* aims interact in deliberation. Here, I refer to the work of Steglich-Petersen (2009), in which he distinguishes between how *mutually compatible* and *mutually incompatible* aims interact. In the first sense, mutually compatible aims interact in such a way as *not* to exclude the pursuit of each other. Aiming to watch a movie and to go to the gym are again clear examples. In ordinary circumstances, when you weigh these aims in deliberation, you can decide on a course of action that allows you to pursue both (and, if all goes well, to satisfy both). You do not need to abandon one aim in favour of the other—they are mutually compatible. This kind of interaction between aims is possible among a whole variety of aims that we pursue on a daily basis: many of the actions we perform are consistent with the pursuit of other aims. In the second sense, mutually incompatible aims *do* exclude the pursuit of each other. Steglich-Petersen gives the following example: imagine that you want to spend the summer perfecting your golf swing, but that you also want to catch up on your research duties. Given that playing golf for the whole summer is going to prevent you from making sufficient progress on your research (and vice versa), you cannot decide to do both. In deliberation, you must choose the aim that you prefer and abandon the other. Any attempt to carry out both aims will only undermine the possibility of achieving either. This is also a kind of interaction that we are familiar with: often, we find that we must abandon certain aims in favour of others.⁶

⁶ As Steglich-Petersen acknowledges, the point is not that mutually incompatible aims logically or necessarily exclude each other (though they might); it is just that they practically exclude each other for the individuals who are weighing them in deliberation. Given that we are interested in how aims interact in actual cases of deliberation, this weaker (practical) modality is sufficient to demonstrate that deliberation sometimes requires us to abandon one (or more) aims in favour of other aims.

Thus, there are, broadly speaking, two kinds of interaction that can take place between aims. However, despite their differences, they are similar insofar as they both involve a process of instrumental reasoning, the overarching aim of which is (ordinarily) to get the most out of the aims under consideration, all things considered. In the mutually compatible sense, this is achieved by the pursuit of multiple aims together (to whatever degree that is possible);⁷ and in the mutually incompatible sense, much the same is true, except that it is achieved only by abandoning one (or more) of our other aims.⁸

With these distinctions drawn, we can outline the following schema for weighing aims in deliberation. In principle, deliberation can take any of the following six structures:

PPC: Compatible interaction between pragmatic aims.

PPI: Incompatible interaction between pragmatic aims.

PEC: Compatible interaction between pragmatic and epistemic aims.

PEI: Incompatible interaction between pragmatic and epistemic aims.

EEC: Compatible interaction between epistemic aims.

E EI: Incompatible interaction between epistemic aims.

The question we are interested in, then, is whether *doxastic* deliberation can take any of these forms. For, if it can, then we can demonstrate that deliberating

⁷ On occasion, pursuing mutually compatible aims decreases the extent to which we can satisfy those aims. For instance, you can watch a movie and go to the gym the same evening; but, as mentioned before, you might have to spend less than your ideal amount of time at the gym, so that you can also get back in time to watch the movie.

⁸ In more detail: The rational thing to do in the golf case is to abandon one of your aims, because attempting to satisfy both will have an overall negative effect on your ability to satisfy either. If you attempt to perfect your golf swing *and* catch up on your research duties, then according to your own reasoning, you will be able to achieve neither of your aims to a satisfactory degree. So, in these situations, you must give one of your aims up completely.

about what to believe does involve multiple aims interacting; and that, therefore, it makes sense to talk about such deliberation in instrumental terms.

To begin our inquiry, we can very quickly rule out any interest in PPC and PPI. These structures are uncontroversial, since we often weigh multiple pragmatic aims in deliberation (as the above examples demonstrate). However, they also do not involve epistemic aims, so they are irrelevant to our concern for belief's aim. Since the aim of belief is said to be the truth, there must be an epistemic aim involved, so we have no reason to analyse PPC and PPI any further. That leaves the last four structures. If doxastic deliberation can take any of these forms, such that multiple aims interact, then the exclusivity objection fails.

In the remainder of this section, I discuss why three of these four structures will not help the teleologists. Specifically, doxastic deliberation does not take the form of PEC, PEI, and EEI. In principle (though not in name), Conor McHugh (2012, 2015) and Asbjørn Steglich-Petersen (2009) have attempted to defend the teleological thesis by appealing to the first two of these structures, respectively. For this reason, much of my attention in the following subsections is on these theorists. Concerning EEI, no serious account has yet been given to suggest that doxastic deliberation can take this form. I thus take the liberty of explaining why such a proposal would be problematic, independently of any serious argument to the contrary. Finally, in Section 2, I argue that doxastic deliberation does take the form of the one remaining structure, EEC.

1.1 PEC AND DOXASTIC DELIBERATION

An inquiry into whether doxastic deliberation can take the form of PEC takes us directly to the core of the exclusivity objection. In fact the problem, as Owens envisions it, can be interpreted as a direct attack on the thesis that the aim of belief can interact with mutually compatible pragmatic aims. As already discussed, the difficulty is that, if we accept exclusivity, then the doxastic deliberator simply cannot take pragmatic considerations into account when forming and abandoning beliefs. In other words, exclusivity rules out the possibility of integrating belief's aim with pragmatic aims.⁹ Thus, for doxastic deliberation to take the form of PEC, we must be willing to abandon the exclusivity thesis.

This is the approach that Conor McHugh (2012, 2015) suggests we should take to defend the teleological thesis.¹⁰ According to McHugh, we must reject exclusivity in favour of a restricted form of pragmatism about belief. In particular, he argues that, on occasion, we can decide to either withhold belief in p or form an outright belief that p in part for pragmatic reasons. This, he claims, is when the evidence we have for p psychologically allows us to believe that p , but nevertheless does not psychologically compel us to believe that p . In other words, we can have psychologically sufficient but non-compelling evidence for believing a proposition. In such cases, we can integrate the

⁹ As Owens acknowledges, there is no principled objection to epistemic aims interacting with pragmatic aims. Contrast *guessing*: when we guess, we aim to guess the truth. This is an epistemic aim. Nonetheless, we can also guess for pragmatic reasons, such as when we guess in order to win a prize on a gameshow. In such situations, we can integrate a pragmatic aim (to win the money) with an epistemic aim (to guess the truth), and take a guess all things considered. Exclusivity denies that belief's aim can interact with pragmatic aims in this sense (see Owens 2003, pp. 290-3 for his discussion of guessing).

¹⁰ McHugh ultimately prefers to attribute an essential knowledge-aim to belief, rather than merely a truth-aim. Nonetheless, he takes the truth-aim to be derivative of the knowledge-aim, so regardless of which aim we prefer, we still need to establish whether we can make any sense of beliefs aiming at the truth. (See McHugh 2011, p. 382).

evidence we have for believing p with our broader pragmatic aims, such that we can either decide to form an outright belief or withhold belief. One example McHugh offers in support of his claim is such that we are faced with strong evidence to believe that a friend is guilty of a terrible crime (there are many reliable witnesses), but we do not want to believe in our friend's guilt since doing so might jeopardise our friendship (especially if he does turn out to be innocent). In such cases, despite evidence being required to form an outright belief (which is why this is a restricted form of pragmatism), the claim is that pragmatic aims can play a role in the doxastic attitude we eventually form: whether that is an outright belief or a withheld belief. Thus, if such cases are telling, then belief's aim can be integrated, under the right conditions, with pragmatic aims. Now, without going too far into the details of McHugh's account (which I have criticised elsewhere),¹¹ there are a couple of points of contention we can raise against this strategy.

First, what makes the exclusivity objection so interesting is how it shows that belief's aim cannot interact with other aims *on the widely accepted assumption* that exclusivity is true. So, while pragmatism would solve the problem, defending the teleological thesis in this way is not going to satisfy theorists who are convinced by the strength of exclusivity. For them, abandoning exclusivity will look a bit like sinking the ship to stop the fire. They would, no doubt, much sooner give up on the teleological thesis than accept pragmatism. Assuming this general preference for exclusivity, then, a more desirable response would be to defend the teleological thesis without abandoning exclusivity.

¹¹ For details, see Atkinson (2018), pp. 43-48.

Second, and no doubt more worrying, are the independent issues that arise if we do accept pragmatism; especially with regard to the wider ambitions of the teleologists. For one, the descriptive ambition of demarcating beliefs from other attitudes according to a truth-aim begins to fall apart. If we can decide not to believe that p , when we have (psychologically) sufficient evidence in support of p , then whatever aim we are pursuing does not look much like a truth-aim. Instead, not to believe that p when we have sufficient evidence supporting p looks like a blatant disregard for the truth. As such, if our descriptive account of belief makes the truth-aim a necessary condition of belief, then the attitudes we form in spite of this aim cannot be beliefs. But this is not something McHugh can accept, since to admit that these attitudes are not beliefs would just be to admit that pragmatism is false.¹² Moreover, it's not so unreasonable to suggest that these attitudes are not, in fact, beliefs. To be sure, they look a lot like what Michael Bratman (1992) calls *acceptances in a context*. Such acceptances can be held and acted on for pragmatic reasons, but there is no reason to consider them beliefs. As Bratman explains, acceptances in a context may conflict with our genuinely held beliefs, such as when we believe that our friend is guilty but act as though he is not *just in case* we are mistaken. So, even in light of the cases McHugh describes, it's not clear why we should accept pragmatism.

In addition, accepting pragmatism causes problems for the teleologist's normative project. Once we agree that pragmatic aims can influence belief, we need an account of correct and rational belief that goes beyond the truth-aim;

¹² For further discussion of the difficulties of accepting both the teleological thesis and pragmatism, see Atkinson (2018).

we need to know when it is correct, if ever, to believe for pragmatic reasons.

Sophie Archer (2015) raises this concern in her criticism of McHugh:

...for the teleologist who affirms exclusivity, the doxastic deliberator... has an overarching aim that has the final say as to whether they are to believe that p —namely, truth. But what of the teleologist who denies exclusivity, such as McHugh? (Archer 2015, p. 10)

To this question, Archer finds no answer. The problem is that when we accept pragmatism, the normativity of belief takes on a whole new dimension: our account must enter the domain of pragmatic reasons for belief. For the teleologist committed to exclusivity, there is no such problem. They have one, and only one, requirement: to believe the truth. But for the pragmatist things are not this simple. They must account for belief rationality in a way that is additionally sensitive to pragmatic reasons, and it's not at all clear how this can be done.¹³

For these reasons, appealing to pragmatism is an uncomfortable route for the teleologist to take to avoid the exclusivity objection. There is, at least, considerably more work to be done. At this point, then, we can conclude that doxastic deliberation does not take the form of PEC. This leaves us with three more options to explore.

1.2 PEI AND DOXASTIC DELIBERATION

The next structure on our schema is PEI. In doxastic deliberation, can incompatible pragmatic and epistemic aims interact? As with PEC, this

¹³ For additional criticisms of McHugh's position, see Sullivan-Bissett (2017) and Sylvan (2016).

structure involves pragmatic aims, but it does not, as we shall see, require us to accept pragmatism. It does not, that is, call for pragmatic reasons to be *integrated* with epistemic reasons for belief. This ensures that we do not automatically run into the same problems from the previous section. Instead, pragmatic aims are said to interact with belief's aim such that the aims involved are in conflict with one another. Steglich-Petersen (2009) defends the teleological thesis along these lines.

According to Steglich-Petersen, belief's aim can interact with mutually incompatible pragmatic aims, in the sense that we can *take up* or *discard* the aim of belief for pragmatic reasons.¹⁴ For instance, you can decide to take up the aim of belief towards a proposition, because you desire to hold a belief about that proposition (perhaps you desire to hold a belief about when a local restaurant opens, so you begin to look online for the restaurant's opening hours); or you can decide to discard the aim of belief towards a proposition, because you do not desire to hold a belief about that proposition (maybe you are no longer interested in visiting the restaurant, so you stop looking for their opening hours). In each case, pragmatic aims interact with belief's aim, such that they can influence which propositions we decide to take up or discard the aim of belief towards. Notice, then, that none of this requires us to accept pragmatism. Nothing in Steglich-Petersen's suggestion involves beliefs being formed or abandoned *for* pragmatic reasons. Rather, the point is that pragmatic reasons are only relevant to the kinds of beliefs we decide to look into.

¹⁴ Indeed, this is why Steglich-Petersen introduces the distinction between mutually compatible and incompatible aims interacting to the debate.

However, the difficulty with this proposal is that, in the kinds of cases Steglich-Petersen presents, the epistemic aim involved does not look much like the aim of belief. The aim of belief is about considering, in doxastic deliberation, whether to believe that p or to believe that not- p . It is not about considering *whether to look into* whether to believe that p or that not- p . This latter aim looks more like the aim of *inquiry* than the aim of belief. The aim of inquiry is an epistemic aim that can interact with other aims (e.g. by considering whether or not to inquire into the truth of a set of propositions), but it must not be confused with the (quite different) aim of belief. In their criticism of Steglich-Petersen's position, Ema Sullivan-Bissett and Paul Noordhof (2013) make this point explicit. As they explain, in its most plausible form, the aim of belief is:

To believe that p only if p .¹⁵

Truth, it seems, can only be a necessary condition for belief, since if truth were also a sufficient condition, we would, in virtue of being believers, have the quite impossible aim to believe all truths: an aim that we clearly do not possess. Given this formulation, the aim of belief therefore says nothing about what propositions we are required to look into—deciding which set of propositions to look into is a requirement we must put on ourselves, when we deliberate whether or not *to begin an inquiry*. Nonetheless, in Steglich-Petersen's examples, the aim being weighed is precisely *this* aim. For example, when we consider whether to look into the restaurant's opening times, we are not deliberating *what to believe* about those times, but whether or not to inquire into those times at all.

¹⁵ Versions of this aim are defended by McHugh (2011), Sosa (2010), Steglich-Petersen (2006, 2009), and Velleman (2000). This is also the formulation that is the explicit target of the exclusivity objection, as described by Owens (2003).

Carrying out the aim of belief would in fact follow this inquiry, once we had some evidence to consider.¹⁶ For this reason, Steglich-Petersen conflates talk about the aim of belief with another epistemic aim; one that I take to be the aim of inquiry.¹⁷

Thus, the teleologist is once again in trouble. Possible examples of belief's aim interacting with mutually incompatible pragmatic aims do not really, it seems, involve belief's aim. As such, doxastic deliberation does not take the form of PEI. This means that neither of the structures involving belief's aim interacting with pragmatic aims are going to work. We are left, then, with EEI and EEC, and the possibility that *distinct epistemic aims* interact in doxastic deliberation. I now proceed to consider this possibility, which has not yet received the attention it deserves in the literature.

1.3 EEI AND DOXASTIC DELIBERATION

I begin with EEI. The question we must ask is whether mutually incompatible epistemic aims interact in doxastic deliberation. Although pragmatic aims are no longer our concern, the possibility of distinct epistemic aims interacting should still be treated as a live option. But how is this going to work? Unlike in previous sections, I have no explicit theory to reflect on, however I think that there is a principled reason why doxastic deliberation cannot take the form of EEI.

¹⁶ It is perhaps worth pointing out that Owens (2003, p. 288) anticipated this response to his objection, and explained that he didn't think it would work for this very reason.

¹⁷ For their ongoing debate on this matter, see Steglich-Petersen (2009, 2017) and Sullivan-Bissett and Noordhof (2013, 2017).

For this approach to work, we would need to show that there are mutually incompatible epistemic aims for belief's aim to interact with. Given that these aims must be epistemic, the situation would be such that certain epistemic considerations (i.e., those pertaining to epistemic aims that are incompatible with belief's aim) would come into conflict with carrying out belief's aim. We would then need to decide, based wholly on these epistemic considerations, which aim to abandon in favour of the other: belief's aim or the other epistemic aim. Admittedly, this is a bit abstract, however this is because it's difficult to see how there could be such a situation in practice. The issue is that there might not be any epistemic considerations that would ever be incompatible with carrying out the aim of belief. Rather, epistemic considerations of any kind would presumably *inform* us about how to satisfy belief's aim—about whether or not to believe a particular proposition—they would not, it seems, force us to give up that aim.

To see the problem in a bit more detail, consider again that belief's aim is typically expressed as the conditional: to believe that p only if p . Since the aim is a conditional, any change in our epistemic situation would seemingly never require us to abandon it. When we consider whether to believe that p or that not- p , the evidence we have (by our own reckoning) might sufficiently support p . In that case, we can satisfy the conditional by believing that p . But if our epistemic situation changes, such that we no longer have sufficient evidence for p , then we change our attitude accordingly, and either withhold belief in p or believe that not- p —both of which again, under the right conditions, satisfy belief's aim. I can't therefore see how, in doxastic deliberation, we would ever be forced to give up the aim of belief solely for epistemic reasons. In turn, then,

it doesn't look like there could be an epistemic aim in conflict with beliefs aim. Changes in our epistemic situation would not, it seems, require us to abandon beliefs aim, but would only ever serve to guide us when attempting to satisfy it. If this correct, then, doxastic deliberation also does not take the form of EEI.

This leaves our final structure: EEC. Given the failure of earlier suggestions, the teleologists must appeal to EEC in order to defend their position. In the following section, I argue that doxastic deliberation does take the form of EEC, and that therefore the exclusivity objection fails.

2. EEC: DEFENDING THE TELEOLOGICAL THESIS

To see how doxastic deliberation takes the form of EEC, we first need to shift our focus from beliefs aim taken in the abstract, to specific instances (or concrete cases) of doxastic deliberation. Reflect once again on beliefs aim: to believe that p only if p . In the abstract, the aim is to believe only truths. In order to carry out this aim, we consider what epistemic reasons we have for and against believing that p . If we take ourselves to have sufficient reasons in support of p , then we go ahead and believe that p ; if not, then we do not believe that p . This is how the teleologist explains how we go about reasoning in doxastic deliberation. We do so by attempting to satisfy this one particular aim. However—and this is where my argument begins—this is *not the only* aim involved in concrete cases of doxastic deliberation. There is also the aim:

To believe that not- p only if not- p .

In the abstract, of course, this is the same as the previous aim: to believe only truths. However, when we think about these aims (call them *doxastic aims*) in

actual cases of doxastic deliberation, such that the propositions in question are *specified*, they are not the same aim. To demonstrate: imagine that you are considering whether to believe *that it is sunny in Darley Dale* (only if it is...) or to believe *that it is not sunny in Darley Dale* (only if it is not...). When the propositions are specified in this way, the two aims are quite clearly distinct: aiming to believe a proposition (only if it is true) is not the same as aiming to believe the negation of that proposition (only if *that* is true). This is made apparent when we consider that the two aims have different satisfaction conditions. Suppose that it is (in fact) true *that it is sunny in Darley Dale*. If doxastic deliberation involves only the first aim, then you do not violate it by believing *that it is not sunny in Darley Dale*. (In other words: you do not violate the aim to believe p only if p by believing not- p when p is true.) However, under the same conditions, you *do* violate the second aim: according to the second aim, you cannot believe *that it is not sunny in Darley Dale* when, actually, it is. (This time: you violate the aim to believe that not- p only if not- p by believing not- p when p is true.) This is the point that the teleologist must acknowledge in order to understand how *weighing* takes place in doxastic deliberation, and is thus the position they must adopt in order to avoid the exclusivity objection. They must agree that there are (at least) these two doxastic aims interacting in any specific instance of doxastic deliberation.¹⁸ Deliberating about what to believe involves a careful negotiation between these two aims, which is a matter of considering the epistemic reasons we have in support of a proposition and its negation, such that the beliefs we then hold satisfy both.¹⁹ If this is correct, then doxastic

¹⁸ I say 'at least' because I later argue that there are potentially many aims interacting in doxastic deliberation, insofar as we might be deliberating over multiple propositions.

¹⁹ One advantage of my proposal, besides helping us to avoid the exclusivity objection, is that it explains why we should not (and perhaps cannot) believe that p when not- p is true (and vice versa). Believing falsehoods is always a violation of one of our doxastic aims. However, a difficulty worth mentioning, which Sullivan-Bissett (2018) raises against both teleological (and normative) theories of belief, concerns

deliberation takes the form of EEC. The doxastic aims involved are both epistemic, since they require us to believe only truths, and they are compatible, since neither must be discarded in favour of the other. This, I take it, is true for all instances of doxastic deliberation.

At this point, it is helpful to reflect for a moment on what I take to be the significance of my observation, especially in response to a potential objection. It might be noted that by continuing to accept exclusivity, an important difference remains between my proposed doxastic aims and other non-doxastic aims. Reasons for belief, unlike reasons that motivate other aims, are restricted to one modality—namely, they are exclusively epistemic. This is the essence of the exclusivity thesis which I do not wish to deny. However, other aims, including other epistemic aims, can interact across modalities. As an example, take the aim of inquiry, as discussed in Section 1.2. The aim of inquiry is (something like) to uncover certain truths about a particular set of propositions (that is, about a particular topic). This means that epistemic reasons are essential for satisfying this aim. Nonetheless, pragmatic aims are also important in deliberation about what to inquire into; for, as mentioned in Section 1.2, we

why we should not (and perhaps cannot) withhold belief in p when we have settled on p being true. If the aim (or norm) of belief is to believe that p only if p , then withholding belief is always permissible, regardless of whether we take p to be true. The second doxastic aim I propose—to believe that not- p only if not- p —does not help us to resolve this issue. Even with this second aim, withholding is always an option. So, what should we make of withheld belief? While I cannot give a detailed analysis of withheld belief here, there are a few possibilities that would be worth exploring. For instance, we might attempt to find a unique aim of withheld belief, such as to give a teleological account of withheld belief that is consistent with our teleological reading of belief (for a related effort, see Sosa 2010, Ch.1). We might also attempt to characterise withheld belief as an altogether different sort of attitude than belief, such that we understand the nature and normativity of withheld belief in and importantly different way from how we understand that of belief—for example, it has been suggested that withheld belief might best be understood as an attitude towards a question, rather than towards a proposition (e.g. see Booth 2014 and Friedman 2013, 2017). Or, alternatively, we might further consider whether Sullivan-Bissett's (2018) functional account of suspension can be squared with our teleological account of belief. On the surface, it is not clear that the functionalists and teleologists are entirely at odds with one another, when it comes to explaining doxastic attitudes. Sure, there are important differences, but they both generally have naturalistic ambitions, and it is common for teleologists to appeal to natural functions when explaining the source of belief's *telos* (e.g. Steglich-Petersen 2006, p. 510, and Velleman 2000, p. 253, fn. 18). Any of these options, if they worked out, could help with the problem of withheld belief without us having to abandon the central claims of the teleologists.

might decide to take up or discard the aim of inquiry for pragmatic reasons. Deliberation about what to inquire into thus permits cross-modality interaction between epistemic and pragmatic aims, and something similar might be said about other aims in general—typically, aims permit cross-modality interaction. For belief, however, this is not the case: exclusivity demands that only epistemic reasons can influence doxastic deliberation. On this basis, it might then be suggested that *this* is the underlying difficulty of the exclusivity objection: the strange feature of belief's aim is *not* that multiple epistemic aims fail to interact in doxastic deliberation, but that these alleged aims do not permit cross-modality interaction with other aims. This is perhaps therefore the issue that the teleologists must address in response to the exclusivity objection. Why is reasoning about belief restricted to one modality when reasoning about other aims is not?²⁰

This, I admit, is not a question that I have attempted to answer. However, as I see it, cross-modality interaction between aims is not what is at stake if we wish to avoid the exclusivity objection. Rather, what the teleologist must demonstrate, in response to the objection, is that there is a role for *instrumental reasoning* in doxastic deliberation. If there is, then describing beliefs as an aim-motivated behaviour makes sense; yet what the exclusivity objection purports to show is that, when a process is necessarily restricted to a singular aim, such that it cannot interact with other aims in deliberation, it doesn't make sense to talk about that "aim" in instrumental terms, so it is meaningless to characterise it as a genuine aim. We might ask: what sense can be made of assessing the means of achieving a particular goal when the assessment of those means

²⁰ Thanks to an anonymous reviewer for urging me to elaborate my stance on this issue.

necessarily cannot affect how the goal is achieved? In such cases, because there cannot be any weighing of the means of achieving the goal (there can be no interaction between aims), it doesn't look like the process can be a truly instrumental one. To give an example, consider the result we get when we attempt to characterise a (more obviously) non-aim-motivated process according to a singular aim. Take the "aim of hair" which is, say, "to grow". In a sense that parallels the exclusivity objection, the difficulty with interpreting this "aim" as a genuine aim is that, since it cannot interact with other aims in deliberation, the process cannot be explained in instrumental terms. We don't sit around deliberating over the best way to carry out the "aim of hair" by weighing this "aim" with other aims and thereby decide whether or not (or how) to grow hair—rather, grows regardless of any other aims that we might have. In this sense, any attempt to characterise hair growth in instrumental terms is implausible, because it is impossible to weigh the "aim of hair" with other aims in deliberation. In turn, this means that such talk of the "aim of hair" cannot be literal, and must at best be thought of as a metaphorical description of what hair is doing (if we so wished to describe hair in such terms). However, as we should now be able to see, believing is not like hair growth. When we deliberate about what to believe, we do weigh reasons for carrying our various distinct aims, albeit only epistemic ones. We consider how best to satisfy those various aims, and we form and maintain beliefs accordingly. This shows that describing beliefs in instrumental terms makes sense: we form and maintain beliefs by assessing the various means to best satisfy multiple aims. Sure, these aims are wholly epistemic, so there is no cross-modality interaction, but there is interaction nonetheless—and *that* is what is important. When we bear witness to this interaction in practice, which I have attempted to show does take place,

we can see that instrumental reasoning characterises belief; and that, therefore, exclusivity is no obstacle to talking about believing as an aim-motivated behaviour.

Finally, to press the point a little, we can draw an analogy between the aim of belief and (one possible interpretation of) the aim of agency. Imagine, for instance, that we were to accept what Peter Railton (1997) calls the ‘High Brow’ view of agency: that deliberative action necessarily aims at the Good. As Railton points out, this is the practical equivalent of the doxastic claim that beliefs necessarily aim at the truth (p. 62). If we did adopt this view, however, we might wonder whether we are forced to draw the conclusion that agency is not really aim-motivated. On one level, it seems like we have to, since we end up with a practical equivalent of the exclusivity objection. Given that the Good is the *only* aim of agency, it begins to look like there can be no interaction between multiple aims of agency—regardless of what our aims are, we ultimately must aim to do (what we take to be) the Good. In turn, on the assumption that the possibility of cross-modality interaction is necessary between aims, this would mean that agency, like belief, is not really aim-motivated. But this doesn’t seem to be the right conclusion to draw about agency, even on the High Brow view. Instead, we should resist this conclusion by noting that *in specific instances* of practical deliberation, multiple aims interact, but they do so with the Good as their measure of utility. We weigh multiple aims in deliberation, which we think will enable us to maximise utility, where the standard of utility is set by the Good, and we act accordingly. Thus, accepting that there is only one single modality of utility for agency should not force us to say that specific actions are not aim-motivated. My point, then, is just the same for beliefs. At the highest

level of abstraction, truth sets the measure of utility for belief; but in specific instances of doxastic deliberation distinct epistemic aims are weighed in pursuit of the truth. Therefore, nothing about restricting an activity's measure of utility to a single modality undermines the fact that multiple aims can interact in order to satisfy that standard.²¹

That brings us to the end of my general defence of the teleological thesis. To summarise, my response severs the connection between P2 and P3 of the exclusivity objection, as it is outlined in the Introduction. When we focus on specific instances of doxastic deliberation, we bear witness to (at least) these two epistemic aims interacting: to believe that p only if p , and to believe that not- p only if not- p . For that reason, the exclusivity thesis does not undermine the fact that multiple distinct aims interact in doxastic deliberation. This shows that reasoning about what to believe can be interpreted as an instrumental process, and that the aims involved are, therefore, genuine aims.

2.1 WEIGHING MORE EPISTEMIC AIMS IN DOXASTIC DELIBERATION

At this final juncture, the issue might be pressed that, even if there are, as I have said, two doxastic aims involved in doxastic deliberation, this, at best, only just meets the requirement that we have set on aims interacting; so I have not done

²¹ A potentially important disanalogy between agency and belief, even on the High Brow view, is that we can be practically akratic but we cannot be doxastically akratic. That is, we can recognise that x is Good, but not be compelled to do x ; yet we cannot recognise that p is true, but not be compelled to believe that p . Thus, in this sense, there is still an important difference between practical and doxastic deliberation. However, while I think this distinction is probably true, I do not think I need to deny practical akrasia to make my point. Rather, the strength of my comparison holds on the basis of practical akrasia being hypothetically impossible: *even if* practical akrasia were impossible, such that we could only be motivated by the Good, I still think we would not be forced to conclude that agency is not an aim-motivated activity. It would still make sense to talk about specific actions as aim-motivated, even though the only truly motivating reasons for those actions would be reasons that the agent took to be in service of the Good.

enough to fully ease the pressure of the exclusivity objection. The fact that there are only two doxastic aims is almost as peculiar as the presence of only one. What kind of aims can only interact with one other aim? Practical deliberation, in contrast, allows for the possibility of *many* distinct aims interacting. When we consider what actions to pursue, there can be many options—to do action A, B, C, D, etc. In fact, there is seemingly no logical limit to the number of aims that can be involved in an instance of practical deliberation (although there must be a psychological limit). So, why would doxastic aims be so limited to only two?

To this concern, my response is to demonstrate—in an extension of my original proposal—that, in any given instance of doxastic deliberation, there can, in fact, be *many* aims interacting. The first assumption to bear in mind, which I used to support my argument in the previous section, is that aims can be distinguished according to their satisfaction conditions. In this respect, having distinct satisfaction conditions is sufficient for distinguishing between aims. With this assumption in mind, we can see that we cannot *only* distinguish between *aiming to believe that p only if p* and *aiming to believe that not-p only if not-p*, but also between many other doxastic aims that involve believing distinct propositions. For instance, consider the following case, modified from Matthias Steup (2017, p. 2681): You believe p , q , and $(p \wedge q) \rightarrow r$, but then you get strong evidence for not- r . In this situation, if you do not revise your beliefs, you will end up with an inconsistent belief-set. Thus, to make sure your beliefs are consistent, you have the option to reconsider the beliefs you hold in a number of ways. While holding the belief that not- r , you can decide to drop the belief that p , drop the belief that q , drop the belief that $(p \wedge q) \rightarrow r$, or drop any combination of these

beliefs; alternatively, you can reassess your evidence for not- r , and instead drop the belief that not- r . Each of these options enables you to maintain a consistent set of beliefs. Now, the point I am making is that, on my account, we can interpret such instances of deliberation as involving many distinct epistemic aims interacting. As we have said, we can distinguish aims according to their satisfaction conditions, and the different propositions involved set precisely that: distinct satisfaction conditions for believing. For instance, *aiming to believe that p only if p* and *aiming to believe that r only if r* are not the same aims when p and r are specified. This is apparent if, for example, we specify ' p ' as 'today is Sunday' and ' r ' as 'the lawn needs mowing'. Whether these aims are satisfied would depend on completely different states of affair obtaining, and thus deliberating over which beliefs to hold would involve different aims—in just the same way as deliberating whether to believe the negations of those propositions would involve different aims. Thus, on this basis, we end up with the potential for many doxastic aims interacting in doxastic deliberation. In the above example, taken from Steup, we can observe:

The aim to believe that...

... p only if p

... q only if q

...not- r only if not- r

... $(p \wedge q) \rightarrow r$ only if $(p \wedge q) \rightarrow r$

...and so on, for the various combinations of these aims.

So, while we can abstract away from any given instance of reasoning about what to believe, by pointing out that each of these aims require us to believe only

truths, it remains the case that when the propositions are specified, the aims involved are not the same. For this reason, we can go so far as to say that the number of aims involved in an instance of doxastic deliberation depends, in part, on the number of propositions under consideration.²² This final point again reveals that, while practical and doxastic deliberation are not exactly the same, they do have important parallels. Both have the potential to involve a multitude of aims interacting, such that the aims are clearly distinguishable on the basis of their different satisfaction conditions.²³ This, I hope, should be enough to alleviate the worry that doxastic deliberation is peculiar insofar as it only involves two aims; since, despite my argument in Section 2, doxastic deliberation need not only involve two epistemic aims, but may involve many.²⁴

CONCLUSION

The exclusivity objection takes the exclusive role of epistemic reasons in doxastic deliberation to be an obstacle for the teleological thesis. Since belief's aim does not interact with other aims, it is claimed that the aim of belief is not much like an aim at all. However, while this problem holds across a range of structures for weighing aims—namely PEC, PEI, and EEI—it does not hold for EEC. I have endeavoured to show that, despite exclusivity, distinct mutually compatible epistemic aims do interact when we consider what to believe. This can be observed when we focus on specific instances of doxastic deliberation,

²² This would not be a one-to-one correspondence, since the number of aims would also depend in part on the relations that are thought (by the agent) to hold between the propositions in question.

²³ Matthias Steup's *doxastic decision thesis* is pertinent to the comparison I am making. According to this thesis, when we deliberate what to believe, we subsequently carry out *intentions* to believe the propositions that we take ourselves to have the best reasons to believe, just as we carry out intentions to perform actions (Steup 2012, p. 158; also see his 2017). My point in the text can also be put in terms of intentions to believe: if we understand doxastic aims as intentions, then it makes sense to describe doxastic deliberation as a process that involves a multitude of intentions to potentially be carried out, where those intentions differ according to the propositions that we intend to believe.

²⁴ Thanks to an anonymous reviewer for this journal, who encouraged me to be clear on this problem.

such that the propositions being deliberated over are specified. If this is correct, then we are again free to give a teleological account of belief, one that I hope is consistent with the teleologists' ambitions of accounting for the descriptive and normative features of belief in doxastic deliberation.^{25, 26}

²⁵ I limit my success to the domain of doxastic deliberation because that it is the focus of the exclusivity objection. There are, of course, other objections that the teleologist must face. One particularly pressing concern is Shah's (2003) 'teleologist's dilemma', which I do not profess to have solved here. Nonetheless, I take my proposal to be consistent with potential solutions to the dilemma, such as that proposed by Steglich-Petersen (2006).

²⁶ **Acknowledgments:** First of all, I'd like to thank my Ph.D. supervisors, Darrell Rowbottom and Jenny Nado. It is due to their guidance during the Ph.D. process, and their many helpful comments on drafts of my thesis, that I have been able to write this paper. I'd also like to thank Andrea Sauchelli, Derek Baker, and Tinghao Wang for their detailed comments on earlier drafts of this paper, and an audience at the 2018 *Frontiers of Epistemology* conference at Yonsei University, Seoul, for their helpful feedback. And finally, I'd like to thank two anonymous referees for this journal, whose insightful comments helped me to clarify many important points.

REFERENCES

- Archer, Sophie. (2015). Defending Exclusivity. *Philosophy and Phenomenological Research* 93 (2).
- Atkinson, C. J. (2018). *On Belief: Aims, Norms, and Functions*. Lingnan University: Ph.D. Dissertation.
- Booth, Anthony Robert. 2014. 'Two Reasons Why Epistemic Reasons Are Not Object Given Reasons'. *Philosophy and Phenomenological Research* 89 (1): 1–14.
- Bratman, Micheal (1992). Practical Reasoning and Acceptance in a Context. *Mind* 101 (401):1-16
- Chan, Timothy (ed.) (2013). *The Aim of Belief*. Oxford University Press.
- Friedman, Jane (2013). Suspended Judgment. *Philosophical Studies* 162 (2): 165–181.
- (2017) Why Suspend Judging? *Noûs* 51 (2): 302–326.
- Kelly, Thomas (2003). Epistemic rationality as instrumental rationality: A critique. *Philosophy and Phenomenological Research* 66 (3):612–640.
- McHugh, Conor (2011). What Do We Aim At When We Believe? *Dialectica* 65 (3):369-392.
- (2012). Belief and aims. *Philosophical Studies* 160 (3):425-439.
 - (2015). The Illusion of Exclusivity. *European Journal of Philosophy* 23 (4):1117-1136.
- Owens, David J. (2003). Does belief have an aim? *Philosophical Studies* 115 (3):283-305
- Railton, Peter (1997). On the hypothetical and non-hypothetical in reasoning about belief and action. In Garrett Cullity & Berys Nigel Gaut (eds.), *Ethics and Practical Reason*. OUP. pp. 53-79.
- Shah, Nishi (2003). How truth governs belief. *Philosophical Review* 112 (4):447-482.
- Shah, Nishi & David Velleman, J. (2005). Doxastic deliberation. *Philosophical Review* 114 (4):497-534.
- Sosa, Ernest (2010). *Knowing Full Well*. Princeton University Press.
- Steglich-Petersen, Asbjørn (2006). No Norm needed: On the aim of belief. *Philosophical Quarterly* 56 (225):499–516.

- (2008). Does doxastic transparency support evidentialism? *Dialectica* 62 (4):541-547.
- (2009). Weighing the aim of belief. *Philosophical Studies* 145 (3):395-405.
- (2017). Weighing the Aim of Belief Again. *Logos and Episteme* 8 (1):141-145.

Steup, Matthias (2012). Belief control and intentionality. *Synthese* 188 (2): 145-163.

- (2017). Believing Intentionally. *Synthese* 194 (8): 2673-2694.

Sullivan-Bissett, Ema (2017). Aims and Exclusivity. *European Journal of Philosophy* 25 (3):721-731.

- (2018). Explaining doxastic transparency: aim, norm, or function? *Synthese* 195 (8):3453-3476.

Sullivan-Bissett, Ema & Noordhof, Paul (2013). A defence of Owens' exclusivity objection to beliefs having aims. *Philosophical Studies* 163 (2): 453-457.

- (2017). Another Defence of Owen's Exclusivity Objection to Beliefs Having Aims. *Logos and Episteme* 8 (1): 147-153.

Sylvan, Kurt (2016). The illusion of discretion. *Synthese* 193 (6):1635-1665.

Velleman, David (2000). On the aim of belief. In *The Possibility of Practical Reason*. Oxford University Press. pp. 244--81.