

This is the preprint version of an article that has been accepted for publication in *The Philosophical Quarterly* published by Oxford University Press. Please do not cite without permission of the authors.

Epistemic Judgment and Motivation

Abstract: Is there an epistemic analogue of moral motivational internalism? The answer to this question has implications for our understanding of the nature of epistemic normativity. For example, some philosophers have argued from claims that epistemic judgment is not necessarily motivating to the view that epistemic judgment is not normative. This paper examines the options for spelling out an epistemic analogue of moral motivational internalism. It is argued that the most promising approach connects epistemic judgments to doxastic dispositions, which are related to motivation in a fairly tenuous sense. It is also argued that this approach currently lacks a plausible and informative account of the nature and workings of these doxastic dispositions, and, hence, an explanation of the range of phenomena internalist theses typically set out to explain. The most promising route for developing such an account, based on recent expressivist work, is investigated and found inadequate for the task.

1. Introduction

The connection between moral judgment and motivation is strikingly systematic and reliable. A core question in meta-ethical theorizing is: what explains this connection? One explanation is that moral judgments are in some way *necessarily* connected to motivation. Perhaps moral judgment itself gives one a motivating reason to Φ , or perhaps it is otherwise necessarily connected to being motivated to Φ (e.g. by being necessarily connected to a distinct conative state). These are versions of moral motivational internalism:

Moral Motivational Internalism ('MMI'): Necessarily, by way of an internal connection, whenever an agent A judges that she morally ought to Φ , A is at least somewhat motivated to Φ .¹

The basic idea is to explain the connection between moral judgment and motivation in terms of moral judgments' *nature*. Motivational externalists reject MMI. They claim that whatever

¹ It is well-known that MMI must be qualified in various ways. For example, whatever claim can plausibly be made about a necessary connection between moral judgment and motivation, it is surely a *pro tanto* claim. Furthermore, there is plausibly a necessary connection between moral judgments and motivation only so long as the relevant person meets certain conditions (such as being 'psychologically normal', being (practically) rational, etc.). All of these issues, and potentially more, will probably also arise in the epistemic case. We will mostly bracket them for the sake of focusing on the broader question, but the reader should bear in mind that all discussed claims will have to be suitably qualified.

connection there is between moral judgment and motivation, it can be explained by something ‘external’ to moral judgment.

The philosophical interest of the debate between internalists and externalists is in part the upshots these views have for our understanding of the nature of normativity. For example, according to internalists, a key feature of normative judgment per se is its practicality, or tight link with motivation (Korsgaard 1996; Ridge 2014; Smith 1994; Wedgwood 2008). Internalists and externalists largely agree that whatever plausibility MMI has is due to the putatively normative character of moral judgment.

Our starting point in this paper is combining this observation with the observation that there are other kinds of normative judgment besides moral judgment. More specifically, we are interested in the widely held view that *epistemic* judgment—for example, the kind typically expressed by the claims ‘S ought to believe that p’, ‘S knows that p’, or ‘S is rational in believing that p’—is normative. Assuming this, should we expect epistemic judgment to be reliably and systematically connected to motivation, too? Might there be a plausible epistemic analogue of MMI that would explain this – a form of *epistemic* motivational internalism (EMI)?² Answers to these questions would seem to have important upshots for our understanding of the nature of epistemic normativity, as well as our understanding of normativity in general and how different normative domains relate.

For example, Allan Hazlett has recently argued *from* claims that epistemic judgment isn’t necessarily motivating to the view that epistemic judgment isn’t normative (Hazlett 2013, 2014; cf. Côté-Bouchard 2017). This would in turn have deep implications for a wide-range of contemporary debates in epistemology. Consider debates about ‘epistemic norms’—for example, the so-called epistemic norm of assertion or epistemic norm of action. It is common to find these debates couched in normative language, such as ‘permission’ and ‘obligation.’ If it turns out epistemic judgment isn’t normative—or at least not normative in a way that closely parallels moral judgment—we may have reason to re-think the most basic vocabulary of these debates. Somewhat more indirectly, an increasing amount of epistemology focuses on the nature of ‘epistemic blame’ and ‘epistemic blameworthiness’ (Brown forthcoming; Kauppinen 2018; Rettler 2017). The most interesting work explicitly draws on the theory of moral blame and extends it to the epistemic domain. Blame and blameworthiness are closely connected to responsibility, which is in turn a normative notion. If it turns out epistemic judgment isn’t normative—or at least not normative in a way that closely parallels moral judgment—we may have reason to re-think such discussions.

² Archer (2017) asks something similar about the idea that *aesthetic* judgments are normative.

In this paper, we are interested in the very idea that there could be a systematic and reliable connection between epistemic judgment and motivation, one that calls out for explanation in the way meta-ethicists have claimed the connection between moral judgment and motivation does. We think this idea is less straightforward than many seem to believe. *A fortiori*, we think the idea that there is an epistemic analogue of MMI—a claim that some prominent epistemologists have explicitly argued for, and others are implicitly committed to—is less straightforward as well. However, there is one particularly promising way to go, which connects normative judgments—including epistemic judgments—to *dispositions* more broadly, rather than to motivation. We examine this approach in some detail, and argue that more theoretical work is required to make it viable. The aim isn't to argue that such work cannot be successfully done, but to outline the details of a particular kind of explanatory challenge that arises for any attempt to do so.

The paper proceeds as follows. Section 2 imposes some structure on the conceptual terrain and claims that there are essentially three kinds of EMI views worth considering. Sections 3 and 4 examine specific examples of such views. Section 5 turns to a discussion of the relationship between two of them— 'doxastic motivational internalism' and 'dispositional internalism'—and examines the prospects for each of meeting our explanatory challenge. We conclude that, while dispositional internalism is most promising, it stands in need of theoretical work before it can provide a viable account of epistemic judgment and motivation (Section 6).

2. Epistemic Motivational Internalism?

Any form of EMI postulates a relation—'motivation'—between two relata: epistemic judgements and something else. There are different positions to take on the nature of the second relatum, as well as the relation itself. This underpins the options for different forms of EMI.

What does 'motivation' mean in the context of epistemic judgment? We can distinguish two relevant senses. First, there is motivation in the *strict* sense, or 'strict motivation'. Strict motivation is the kind of relation that holds, perhaps most paradigmatically, between desire-like states—desires, intentions, plans, and the like—and the actions they produce. More precisely, strict motivation is the sort of relation holding between a response and a set of mental states that makes the response a fitting target for a folk-psychological *intentional* explanation. If one assumes the orthodox Hume-Davidson model of intentional explanation, strict motivation is the sort of causal relation that holds between responses and sets of desires and beliefs that allow us to properly describe the responses as intentional.

Second, there is motivation in a *loose* sense, or 'loose motivation'. Loose motivation is a broader sort of causal connection between a mental disposition and its output—broader than the

causal connection that holds in those cases in which intentional descriptions are proper. Of course, not just any sort of connection (and, hence, not just any sort of mental disposition) plausibly counts. Even in a broader sense of ‘motivation’ (on which it has a meaning distinct from ‘causation’) not *any* kind of causal relation between a mental disposition and its outputs is a suitable model for the causal operation of normative judgements. For example, when a person is disposed to shake in fear in response to some triggering condition that reminds her of a traumatic event, she *does* have a mental disposition with certain sorts of outputs. But this kind of disposition and how it causally connects to its outputs surely isn’t the kind of connection assumed to hold for normative judgement. Hence, the relevant sort of causal connection that characterizes the kind of disposition connected to normative judgement must be restricted in some way. For example, the connection might be a *rational* connection. We will return later to the issue of how to understand ‘loose motivation.’ For now, all that matters is that loose motivation should be understood as a causal relation between a mental disposition and its outputs that is significantly broader than strict motivation, and doesn’t have to share strict motivation’s characteristic features.

With this distinction in place, what about the second relatum of the relation postulated by EMI? Here, we can distinguish at least two possible positions. First, we might hold that this relatum is some range of *actions* or *intentions*. For example, epistemic judgements might be internally connected to actions of inquiry, such as seeking further evidence, taking another look at the evidence one already has, and so on. Second, we might hold that this relatum is a certain range of doxastic states, such as belief, disbelief, or suspension of belief. For example, judging that you ought to believe that Hartmann is the killer might be connected to coming to believe that he is the killer.

With these distinctions in place, we can distinguish four possible forms of EMI. These fall into two broad camps. In one camp, we have views that postulate an internal connection of *strict* motivation for epistemic judgements. These are forms of ‘Strict Motivational Internalism.’ One possibility here is the view that epistemic judgement strictly motivates a range of intentions and actions. This would seem to be the closest possible epistemic analogue of MMI. After all, a connection of strict motivation to actions/intentions is exactly the sort of connection that proponents of MMI are concerned to identify, albeit between actions/intentions and moral judgments, as opposed to epistemic ones. Let us call this view ‘Traditional Motivational Internalism’ (‘TMI’). Second, there is the view that epistemic judgments strictly motivate a certain range of doxastic states, such as belief, disbelief, or suspension of belief. Such views are still quite close analogues of MMI. After all, they merely differ in their position on *what* the relevant normative judgements motivate, not how they do it. Let us call this view ‘Doxastic Motivational Internalism’ (‘DMI’).

In a second broad camp, we have views that postulate an internal connection of *loose* motivation for epistemic judgements—‘Loose Motivational Internalism’. Such a connection can again either be between epistemic judgments and actions/intentions, or epistemic judgments and doxastic states. But the distinction between these two possibilities won’t matter in what follows. So we will focus on a doxastic states version of the view. We call this kind of view ‘Dispositional Internalism’ (‘DI’). Such a view seems, at first sight, less analogous to MMI than forms of strict motivational internalism—though we will discuss this point further later on. The crucial issue for now is whether *any* of these forms of EMI are plausible.

Table 1.1 Varieties of Epistemic Motivational Internalism

Relatum	Relation with Epistemic Judgment	
	Strict Motivation	Loose Motivation
Actions or intentions	Traditional Motivational Internalism (TMI)	?
Doxastic states	Doxastic Motivational Internalism (DMI)	Dispositional Internalism (DI)

3. Strict Motivational Internalism

In this section we discuss whether and to what extent forms of strict motivational internalism are plausible. We argue that strict motivational internalists face a dilemma, because both TMI and DMI are problematic. On the first horn, we argue that there is a simpler explanation than TMI for the phenomena TMI is supposed to explain. On the second horn, we argue that any form of DMI comes with robust and controversial commitments in the philosophy of mind and action. To make things manageable, we focus our discussion by looking at specific proposals made in the literature.

3.1 Traditional Motivational Internalism

On first sight, it might seem implausible that an internal connection to *actions* and *intentions* is what is characteristic of epistemic judgements. However, as we’ve pointed out above, there may be a connection between epistemic judgments and actions of inquiry. For example, as Kappel and Moeller (2014) observe:

Epistemic Judgment-Motivation Connection (I): there is a reliable and systematic connection between sincere knowledge-attributions of the form ‘S knows that p’ and being motivated to *terminate inquiry into whether p*.

Knowledge-attributions are typically used to express what they call ‘K-judgements’. So according to Kappel and Moeller, K-judgments are reliably and systematically connected to being motivated to terminate inquiry. By inquiry, they mean ‘actions such as taking an extra look at the scene, performing calculations or checking them for errors, going through the steps of a proof, wondering whether the evidence supports a particular conclusion or not, thinking about possible biases that might affect one’s evaluation of the evidence, and so on’ (2014: 1532).

Kappel and Moeller’s proposal is that this observation is best explained by:

K-judgment Internalism: Necessarily, by way of an internal connection, whenever an agent A makes a K-judgment, A is pro tanto motivated to terminate inquiry with respect to p (2014: 1531).

Their aim in defending K-judgment internalism is to deploy it in an argument for epistemic expressivism about knowledge-attributions. According to that argument, utterances of the form ‘S knows that p’ express desire-like states—in particular, desire-like states underpinning an agent’s motivation to terminate inquiry. Here we restrict our attention to K-judgment Internalism.

The view seems to imply that whenever one K-judges, one is *already* engaged in inquiry as to whether p—which is obviously false. But a more nuanced observation can be put as follows. When people make knowledge-attributions of the form ‘S knows that p’, then, if they are engaged in inquiry as to whether p, they are systematically and reliably motivated to terminate inquiry into whether p; if they aren’t engaged in inquiry as to whether p (but acknowledge an interest, and someone suggests inquiring into p, for example) then they are reliably and systematically motivated not to do so (2014: 1532). We’ll just assume for sake of argument that K-judgment Internalism can somehow be modified to explain this.

We might wonder whether the factivity of knowledge attributions, as opposed to anything normative about K-judgments, underlies the observations that K-judgment Internalism is supposed to explain. To judge ‘S knows that p’ is to judge that p, which, in light of considerations about the transparency of belief, arguably just is to believe that p (or at least entails that one believes that p). Perhaps their initial observation really amounts to the observation that there is a reliable and systematic tendency in people not to be motivated to inquire into things they already believe are the case. Such an observation wouldn’t seem to support K-judgment Internalism. But Kappel and Moeller argue that there is no such reliable and systematic tendency in the case of mere belief;

according to them, we must appeal to the normative dimension of K-judgments to account for the sort of behaviour they are interested in.

They appeal to cases in which it is intuitive that a person believes that *p* but would appropriately continue being motivated to inquire into whether *p*. Consider the famous high-stakes bank cases (DeRose 1992; Stanley 2005). Here, a person believes the bank will be open on Saturday, yet it seems appropriate for them to continue being motivated to inquire into whether the bank will be open on Saturday. Indeed, the thought goes, it seems appropriate for them to inquire further precisely because they don't *know* that the bank will be open on Saturday (as described in the case). The upshot is that something beyond the mere factivity of knowledge-attributions underlies the pattern. Perhaps the normative dimension of knowledge-attributions does.

One response is to argue there are similar sorts of cases in which a person intuitively knows that *p*, yet it seems appropriate for them to continue inquiring whether *p*. Brown's surgeon case might be an example (Brown 2008). In general, assuming knowledge doesn't require credence 1, there can be cases in which an attributor judges that *S* knows that *p*, but remains unsure whether *p*. When the stakes are high enough, it may be natural to expect the attributor to continuing inquiring whether *p*.³ These are controversial issues, and addressing them in detail would take us too far afield. In any case, we are unconvinced by K-judgment Internalism for other reasons.⁴

In particular, it's not obvious that K-judgment Internalism is the best explanation of the connection between K-judgment and inquiry. For example, an alternative explanation is that knowledge is the goal of inquiry. Call this the K-goal thesis:

K-goal thesis: For any agent *A*, *A*'s inquiry into whether *p* is successful, from the epistemic point of view, when and only when *A* comes to *know* that *p*.

The basic idea is that agents inquiring whether *p* who come to make a K-judgement, naturally *lose* their motivation to inquire into whether *p*, and thereby become—as Kappel and Moeller would put it—'motivated to terminate inquiry' into whether *p*. But, this is just an instance of the general phenomenon that agents tend to lose their motivation to achieve a goal when they judge themselves to have accomplished the goal. When these factors obtain, they combine to constitute an agent's motivating reason to terminate inquiry into *p*. Thus, this explanation is externalist insofar as it requires appealing to motivational states not internal to K-judgment itself.

³ Thanks to [omitted] for raising this point.

⁴ Kappel and Moeller consider and discuss the surgeon case objection (2014: 1538), but do not address the general recipe proposed here for constructing counterexamples to K-judgment Internalism.

One might challenge the idea that knowledge is *the* goal of inquiry (Kappel and Moeller 2014: 1547). Perhaps there are many such goals—for example, perhaps sometimes it is having a *sufficient reason* to believe, and perhaps sometimes it is *certainty*. However, our point only requires that knowledge is *a* goal of inquiry, which it very plausibly is. What is doing the explanatory work in this appeal to the K-goal is simply that people’s tendency to lose their motivation to achieve a goal when they judge themselves to have achieved it is a general phenomenon. Appeal to the K-goal has explanatory power simply insofar as it would subsume a particular case under this more general phenomenon. Even if there are other goals of inquiry, it’s quite plausible that knowledge is one of them, and that whenever people judge themselves to know things, they tend to stop inquiring because they’ve met this particular goal of inquiry (or some goal it encompasses).

Why think knowledge is *a* goal of inquiry? Kappel and Moeller may be able to defend K-judgment Internalism at this level of the dialectic. Recall, their aim in defending K-judgment Internalism is defending a form of epistemic expressivism. They might agree that Epistemic Judgment-Motivation Connection (I) is best explained in terms of the idea that knowledge is a goal of inquiry. But they could insist that this latter idea would in turn be best explained by the idea that utterances of the form ‘S knows that p’ express certain desire-like states.⁵ Such explanatory connections may lend support to K-judgment Internalism after all. But we remain pessimistic. For a start, this approach comes with the issues for K-judgment Internalism discussed above. It comes with worries about the factivity of knowledge attributions, and the general recipe for counterexamples to K-judgment Internalism. While those issues remain far from conclusive, there are other attractive explanations of the assumption that knowledge is a goal of inquiry. For example, perhaps knowledge is a goal of inquiry because—as Fantl and McGrath (2009) have argued—knowledge that p is necessary and sufficient for treating p as a reason for further action and belief. On this approach, knowledge is a highly valuable epistemic state; as such it is a natural goal of inquiry.

On a more general note, although we shouldn’t expect EMI to be exactly like MMI, K-judgment Internalism seems disanalogous to MMI in the wrong sort of way. The pattern being explained here seems open to explanation in terms of considerations about the goal of inquiry, rather than considerations about the nature of epistemic judgment. These aren’t differences we should expect to find simply in light of the fact that we are comparing moral judgment with epistemic judgment.

⁵ Thanks to [omitted] for raising this point.

While we've focused specifically on Kappel and Moeller's proposal, the lessons learned seem to apply to any form of TMI. TMI postulates a relation between epistemic judgment and actions and it is simply difficult to imagine what other sort of actions epistemic judgment could be internally connected to besides inquiry.⁶ Of course, we may be missing some important conceptual possibility. If so, whatever that view is would be worth exploring.

3.2 Doxastic Motivational Internalism

Perhaps there is a reliable and systematic connection between epistemic judgment and motivation to *believe*, or be in some other doxastic state. Veli Mitova (2011) is a good example of someone who thinks so. She makes the following claim:

Epistemic Judgment-Motivation Connection (II): 'It seems trivial [...] that one cannot be genuinely curious about some aspect of the world, and not be moved by at least a class of one's epistemic judgments.' (Mitova 2011: 60).

She argues this is best explained by:

Ought-to-believe Internalism: There is a necessary connection between making an epistemic judgment and being moved to believe in accordance with it (Mitova 2011: 58).

For our purposes, what is significant about this sort of view is that, rather than positing some kind of internal connection between epistemic judgment and action, it posits a direct connection between epistemic judgment and *belief*. Furthermore, Mitova herself explains what she means by 'being moved' as follows: 'being motivated, or moved, by a normative judgment means that one either tries to comply with it, or recognizes that one has done something wrong in not trying to comply with it' (2011: 59). In our terms, she seems to have strict motivation in mind.

⁶ This would explain why Kappel and Moeller's view is the only form of TMI we've come across in the literature. Recently, Michael Ridge (2018) has challenged Kappel and Moeller on the basis of the worry about transparency of belief considered above. In his discussion, he goes on to develop a sophisticated view about the relationship between epistemic judgment and motivation. One way of developing the view is to take a line similar to Kappel and Moeller's, combined with a thesis about *pragmatic encroachment*, rather than direct claims about knowledge-ascriptions (Ridge 2018: 154). Ridge's proposal requires endorsing pragmatic encroachment of a fairly robust variety (a stakes based, SSI-style encroachment). This is of course controversial, but further discussion would take us too far afield, so we will simply leave this as a choice point in the dialectic.

Interestingly, there are meta-normative theorists who don't directly engage with the issue of whether there is an epistemic analogue of MMI, but who are nevertheless committed to the idea, and who moreover seem to be committed to strict motivation in a way similar to Mitova. Consider the views of those who have defended epistemic expressivism in other ways than by starting from claims about motivational internalism. Expressivists hold that normative judgments *are* (at least partially) desire-like states. For example, Allan Gibbard (2003) suggests that epistemic judgments (in particular knowledge-attributions) are *plans* to rely on the knower's judgment in one's own belief formation. He (2012) is also explicitly committed to the idea that there are plans to *believe* (or at least to have certain sorts of credences). And Michael Ridge (2007) suggests that epistemic judgments partially consist in 'plans about how to go about revising and updating our beliefs, and that these plans can, perhaps indirectly, have a bearing on what beliefs we come to abandon or adopt' (2007: 101) and involve 'epistemic endorsement of certain procedures for deciding what to believe' (2007: 103). These theorists, too, seem committed to the idea that we can be strictly motivated to believe.

While DMI may avoid the worries faced by TMI, it raises other issues. First, how often do people believe things on the basis of forming *epistemic* judgments? Typically, people seem to believe the things they do, for example, because of observations they make about the world, or things they hear from others—not because they make overt epistemic judgments. Second, if people form beliefs simply upon judging that *p*, we might wonder what makes Epistemic Judgment-Motivation Connection (II) special. In what way does it highlight anything particularly interesting about epistemic judgment, *per se*? These are important questions. But at least two points are immediately relevant in response.

First, there are more complex kinds of cases in which people do make overt epistemic judgments and seem to adopt certain doxastic attitudes in response. For example, you are conducting a scientific experiment and thinking about the evidence. You find yourself saying, 'I wish the results had been different, but the evidence supports *p*, so I believe that *p*', and thereby coming to believe that *p*.

The second point is the following. Even when people believe things on the basis of perception, for example, they are plausibly guided (however implicitly) by background epistemic judgments to the effect that beliefs should be formed on the basis of sufficient evidence (for example), and about what counts as evidence for what, and so on. The reason *why* people form beliefs about the environment in the way they do is plausibly partially explained by such background regulative epistemic judgments. So, if the reader is concerned about the idea of a thesis that focuses on seemingly rare cases of agents forming overt epistemic judgments, we invite them

to shift their focus to these regulative judgments. The relevant point remains the same.

Perhaps most importantly, though, this form of internalism seems to imply that belief is intentional. More specifically, the suggestion that people can be (strictly) motivated to believe seems to commit one to the idea that people can believe at will. This is the controversial thesis of doxastic voluntarism. Orthodoxy would have it that (at least as a matter of psychological fact) the extent to which people can believe things at will—in a direct way, such as forming the belief right now that you’re not reading an English sentence—is extremely limited, if possible at all (Alston 1988). Some philosophers even argue that believing at will is conceptually impossible (Williams 1973; Adler 2002; Church 2002; Buckareff 2014).

Of course, there is a lot to say here. For a start, Mitova has her own responses to worries about doxastic voluntarism.⁷ She makes two relevant points. First, she claims that action itself is no more voluntary than belief ‘in a totally brute unconstrained sense of voluntariness’ (2011: 57). The idea is that just as we can only believe things for reasons, we can also only *act* for reasons. And in this sense, belief and action are on a par. Second, while we don’t seem to have direct voluntary control over our beliefs, we have various kinds of *indirect* control—such as control over how attentive we are to evidence.

We don’t find these arguments particularly convincing. Regarding the first, Mitova is right to point out that people believe things *for reasons*. But this point has limited utility in the context of defending the kind of voluntarism DMI seems committed to. The kind of voluntarism at issue is to be understood in terms of strict motivation. There seems to be a gap between the idea that belief is responsive to reasons and the idea that epistemic judgment comes with a motivation to believe. The fact that an attitude is responsive to reasons doesn’t imply anything about its relation to motivation. Indeed, it is plausible that there are important differences between the way belief is responsive to reasons and the way our actions are responsive to reasons.

Mitova might respond as follows. She can agree that ‘strict motivation’ should be understood in terms of intentional explanation, but then argue that what counts as an intentional explanation is much broader than we seem to be assuming. In particular, she could argue that intentional explanation simply requires showing that certain states are in some sense responsive to reasons (doing Φ intentionally is just doing Φ for a reason, in the right kind of way). On this picture, belief might seem to count as an intentional activity; and so, by our definition of strict motivation, belief would seem to be something that can be strictly motivated. But then all one is committed to

⁷ Gibbard and Ridge have things to say about voluntarism too. Some of their remarks proceed along lines similar to Mitova’s first point in what follows (Gibbard 2012; Ridge 2018).

in saying that belief can be strictly motivated is the idea that belief is responsive to reasons. And this doesn't seem controversial after all.

This line of thought strikes us as misleading. Even if we concede that doing Φ intentionally is just doing Φ for a reason in the right kind of way, there is a difference between doing Φ intentionally and intending to Φ (Ryan 2003). I might be bending my knees intentionally while I go for an afternoon run, without specifically *intending to bend my knees*. As we understand the idea, a response is a fitting target for intentional explanation if that response is the outcome of a connection between a set of mental states and the agent's intending to Φ (not merely 'intentionally Φ -ing'). Perhaps there is an interesting sense of 'intentionally Φ -ing' which amounts to a kind of responsiveness to reasons. But this is surely very different from what we have in mind when we talk about strict motivation.⁸

Concerning the second argument, many would agree with Mitova that there is some kind of indirect way in which we have control over our beliefs. The problem is that this observation doesn't seem relevant to the project of defending forms of EMI that focus on *doxastic states* against worrying upshots. In particular, the idea that we have (merely) indirect control over our beliefs at most suggests a picture on which our epistemic judgments are connected to certain actions of inquiry (such as being attentive to evidence). Such a connection doesn't call out for anything like DMI as an explanation. Indeed, to the extent that our indirect control goes through actions of inquiry, perhaps this observation calls out for something like TMI as an explanation. But we've already seen that the prospects for TMI as the *best* explanation of the connections we observe between epistemic judgment and actions of inquiry are dim.

Each of these arguments merits an in-depth discussion in its own right.⁹ However, even without doing so there is an important further observation to make. In virtue of being *committed* to doxastic voluntarism, DMI's commitments in the philosophy of mind and action are controversial and robust. After all, doxastic voluntarism isn't compatible with just *any* plausible view in the philosophy of mind and action. So DMI is hostage to the truth of views that are so compatible. Of course, adopting *any* form of motivational internalism is itself a robust theoretical commitment.

⁸ Defending doxastic voluntarism in this way may amount to committing oneself to something more like the claim that belief can be 'loosely motivated'. The possibility that we are misinterpreting certain philosophers as being committed to DMI is something we consider below. Our aim here is primarily to consider DMI as a theoretical option, and outline some of the issues it raises.

⁹ For defenses of doxastic voluntarism, see Ginet 2001; Nickel 2010; Peels 2014. For reasons-responsiveness accounts of voluntary belief, see Owens 2000; Ryan 2003; Heironymy 2009; Steup 2008; Shah 2002.

However, the commitments that come with DMI are far more robust than *just* a commitment to internalism.

Note that things are distinctively different for MMI: nothing in MMI requires us to take on commitments in other areas of philosophy that aren't compatible with any plausible view in these fields. Hence, to adopt DMI is to incur controversial commitments that go well beyond what we are committed to by simply adopting internalism. *Ceteris paribus*, we should prefer views that aren't theoretically costly in this sense. So there is good reason to avoid doxastic voluntarism and, hence, DMI. That is, staying as neutral as possible on doxastic voluntarism is a desideratum for any form of EMI.

This isn't a fatal problem for DMI, but it motivates looking elsewhere. DMI seems to present us with a real epistemic analogue of MMI. But it does so at a significant cost. If there is a less controversial form of EMI, this would, therefore, be an attractive option.

4. Dispositional Internalism

The two forms of Strict Motivational Internalism seem to present a dilemma: on one hand, embracing TMI, it seems right to talk about a judgment-*motivation* connection. But it's far from clear that an internalist thesis is required to explain the connection: when epistemic judgment is connected to intentional states or activities, such as inquiry, there are other motivational states, external to epistemic judgment itself, doing some of the work in explaining the relevant connection. On the other hand, if the proposed connection doesn't concern actions—as in DMI—then, while an internalist explanation of the connection might be plausible, we come up against controversial commitments like doxastic voluntarism. With this dilemma in place, it is worth considering whether DI fares better.

Any form of DI is committed to the idea that we should understand 'motivation' loosely. Bob Beddor (Beddor 2016) has put forward a proposal that seems relevant here. He is interested in defending a form of EMI that focuses on doxastic states—but he frames his thesis in a way that is responsive to the worry that there is something too volitional or intentional-sounding about talk of 'motivation to believe'.

He points out that such talk can be understood as a metaphorical way of putting the more concrete idea that people are disposed to believe things under certain conditions:

Epistemic Judgment-Motivation Connection (III): Typically, if a psychologically normal agent sincerely judges that she's epistemically obligated to adopt doxastic

attitude D towards p, then she'll be at least somewhat disposed to adopt D towards p (Beddor 2016: 54).

Beddor doesn't directly turn to an internalist thesis to explain this. Instead he develops a semantics for the epistemic 'ought' which, if true, would explain this connection. Since Beddor's semantics for the epistemic 'ought' implies an internalist thesis about epistemic 'ought'-judgments, we'll simply take a cue from Beddor and suggest the following:

Dispositional Internalism (DI): Necessarily, by way of an internal connection, whenever an agent A makes a self-directed epistemic obligation-judgment, A is pro tanto disposed to adopt doxastic attitude D towards p.

It should be clear that DI isn't exactly an analogue of MMI. But that doesn't rule out that both DI and MMI naturally fit within a broader unified picture of normative judgements. Perhaps both are particular instances of a general phenomenon. Specifically, it's plausible that all normative judgements come with dispositions to respond in various ways. Perhaps these dispositions differ relative to the respective domains, not merely in what they produce, but also in how they produce it. In the case of the moral domain, for example, we might say these are dispositions to be strictly motivated. In the epistemic domain, the way responses are produced is different. This is one story proponents of DI might tell to explain how their form of epistemic internalism is related to MMI. We said earlier that we shouldn't expect EMI to be exactly like MMI. And according to this story, the differences between DI and MMI are simply differences we should expect to find in light of the fact that we are comparing moral with epistemic judgment.

Note that the difference between DMI and DI isn't merely terminological. It is clear that DMI and DI are substantially different: they involve different claims about the relation between epistemic judgements and belief. However, given the prima facie oddness of 'motivation to believe' talk, one might wonder whether DI is the view that putative DMI theorists really had in mind in the first place. That is, the availability of DI raises the question of whether anyone is actually committed to DMI—or whether there is a reason to be so committed—given that one could equally endorse DI.

Indeed, at this point proponents of DMI—for example, the expressivists above—might

reply that they have been read uncharitably.¹⁰ They might hold that, while their view about the connection between *practical* judgments and motivation is best understood strictly, they can *also* hold the loose view for normative judgements generally. Such an expressivist might hold, for example, that all normative judgments are non-representational states that play a certain role in governing our mental states and events. But this governing role follows the desire-action model only in the case of practically normative judgements. While expressivists such as Gibbard and Ridge explicitly characterize their meta-epistemological views in terms of desire-like states, perhaps we should read notions such as ‘plan,’ ‘endorsement,’ and ‘deciding’ in the epistemic case more loosely. They aren’t talking about plans, endorsements, or decisions in the normal sense, but rather in a technical sense that is more in line with DI.¹¹

This move is definitely possible. But, as we argue next, it is also a move that expressivists, as well as others tempted by DI, should think through carefully.

5. Explanatory Challenge for DI

Despite the potential drawbacks we’ve highlighted, DMI has a significant theoretical benefit that DI doesn’t. First, note that any form of motivational internalism is (at least partially) an attempt to explain something: namely, the way normative judgment is reliably and systematically connected to other sorts of events or states. For example, MMI has been put forward as a thesis that *explains* why and how moral judgment is reliably and systematically connected to action. Why is it so puzzling to judge that one really morally ought to give to charity and yet fail to do so when the occasion arises? It is because of the *internal* connection between moral judgment and action—or

¹⁰ It certainly *seems* as though the authors in section 3.2 are committed to DMI. For example, Mitova appears to explicitly frame her view in terms of strict motivation. And some version of Strict Motivational Internalism seems to fall straight out of the commitments of epistemic expressivists. So it’s not as though we’ve made an obvious mistake in attributing DMI as opposed to DI to those theorists.

¹¹ One point of clarification is important to avoid confusion. Although both Ridge and Gibbard claim that normative judgements are ‘plans’ or attitudes of ‘endorsement’ in a technical sense (see e.g. Ridge 2014: 111-8 or Gibbard 2012: 169-177), the way their notions are technical differs from the way they would have to be technical to count as ways of filling out DI. For example, on Gibbard’s account, plans in his technical sense can do some significant things ordinary plans cannot do. First, while ordinary planning proceeds from one’s own evidence, Gibbardian plans treat this as a special case and allow for planning for all sorts of contingent circumstances where planning proceeds from another person’s evidence (Gibbard 2012: 174). Second, while ordinary planning normally proceeds on the assumption that one will act on the plan, when the occasion arises, Gibbardian plans can be for situations that one *knows* are merely hypothetical (Gibbard 2012: 175). However, Gibbard still seems to hold – and he never says anything to the contrary – that the way plans in his sense govern whatever it is they govern is no different from the way ordinary plans govern intentions (the only difference he seems to hold is a difference in *what* they govern (and what follows from this) not in *how* they govern (Gibbard 2012: 176)). That is, the relation that holds between a plan and a response is supposed to be the same in all cases – the one we find in ordinary plans. And this means that Gibbard’s plans aren’t technical in the way DI requires.

so the internalist explanation goes.

However, while motivational internalism is supposed to explain this phenomenon, the thesis that there is an internal connection between certain judgements and something else, i.e. motivational internalism, is *itself* something we require an account of. What is the nature of this internal relation and how is it that normative judgments bear it to other things? These are questions that plausible internalist theses must answer. For moral judgements, explaining this is part of what Michael Smith (1994) has called ‘the moral problem.’ And, of course, this explanatory challenge arises whenever we postulate an internal connection of the sort characteristic of motivational internalism.

In the moral case, in which the connection of normative judgements to *action* is prominent, we have a clear and obvious folk-psychological model to account for that connection: the reliable and systematic connection between desire-like states and action that has been at center stage in the meta-ethical debate surrounding motivational internalism. Specifically, wherever we assume that the connection is *strict motivation*, we have a clear folk-psychological model to account for the connection. In fact, this is exactly how most authors account for MMI, even those who aren’t expressivists (see e.g. Copp 2001; Smith 1994; Tresan 2006).

DMI shares this advantage with MMI: it explains the reliable and systematic connection between epistemic judgments and other responses *and* also provides an informative model of the nature of this internal connection in terms of an independently understood folk-psychological model. To give an account of the relation between epistemic judgements and beliefs, proponents of DMI only need to appeal to strict motivation, which is something we have a robust independent understanding of. Of course, the applicability of this model in the case of DMI is controversial, but this doesn’t change the fact that proponents of DMI *have* a model for the internal connection they postulate, as well as models as to why this connection holds.

On first sight, the situation is quite different for forms of DI, which don’t seem to offer an equivalent model to account for the connection. In fact, the formulation of DI we have considered here on behalf of Beddor seems to merely reformulate the claim that there is an internal connection between epistemic judgements and beliefs in terms of dispositions. Clearly, this isn’t really an *account* of that connection: it tells us neither what the relation between epistemic judgement and belief is, nor does it explain how epistemic judgement comes to stand in that relation to beliefs. It merely tells us *that there is a causal relation* between epistemic judgement and belief such that when one forms the former one also forms the latter. To make the account fully feasible, more needs to be said. This brings us back to the notion of ‘loose motivation’ we introduced earlier. What this challenge essentially amounts to is to give an account of what ‘loose motivation’ is, as well as

an account of in virtue of what loose motivation holds between two relata.

This is a serious theoretical challenge and a response must meet several desiderata. First, it must plausibly distinguish loose motivation from other causal connections, and must give an account of the kind of disposition that is relevant for loose motivation. Second, loose and strict motivation must be clearly distinguishable. Specifically, loose motivation shouldn't exemplify the features of strict motivation that raise worries for DMI. But, third, either strict motivation must be a *special* case of loose motivation or both loose and strict motivation must be special cases of a more general connection. This requirement derives from the fact that both MMI and EMI are to be seen as instances of a more general phenomenon that characterizes *all* normative judgements. Specifically, if some connection to *motivation* is taken to be what is distinctive of *normative* judgements – the guiding assumption of our inquiry here –, the way in which normative judgements (strictly and loosely) motivate must, plausibly, be related in some way. Hence, anyone who endorses DI faces a significant theoretical challenge, namely to provide a general and unified theoretical account of the way normative judgment produces other sorts of responses. Meeting these desiderata is, by no means, trivial. In fact, on first sight it isn't exactly clear whether folk-psychology has the resources for such an account — one that is truly general and can properly account for the individual cases.

Let us end by raising the stakes and criticizing what is a natural way of trying to meet this challenge. This is to build an account of loose motivation by using *strict* motivation as a model while trying to abstract away from the features of strict motivation that are problematic in, for example, the epistemic case. Given we have a good and independently understood model for strict motivation, this seems like a good starting point for meeting our challenge.

To illustrate how this idea might be pursued, consider the following suggestion that focuses on *plans* as the basic model.¹² Consider plans in the ordinary sense and how they produce actions. What is distinctive about plans is that they play a structuring, organizing, regulating, and controlling role vis-à-vis our intentions. That is, the way in which *plans* strictly motivate actions is in virtue of the *regulative role* they play in our mental economy with regards to intentions. Once we have this role in view, we might hold that mental states are able to play such a role not (just) for intentions but (also) for other mental states, such as beliefs or emotions. If this is correct, abstracting away from the specific role of plans in the ordinary sense to a more general phenomenon – namely a regulative *plan-like* role vis-à-vis mental responses –, we now have a model for thinking about

¹² There are other states one might use as the basic model. We take it, though, that what we argue here applies *mutatis mutandis* to all cases in which the basic model are states that strictly motivate, i.e. where the broader relation is explained in terms of the narrower one. Thanks for [omitted] for drawing us out here.

mental states with a particular feature, namely a certain regulating role in the formation of other mental states. Note, though, that such mental states – ‘plans’ in a more technical sense of the word – would then produce responses in a way that is *similar* to (a certain form of) strict motivation, but which isn’t quite strict motivation. After all, such mental states would operate exactly as plans in the ordinary sense do, except that they wouldn’t regulate intentions. If the way that plans strictly motivate, however, has something to do with their role vis-à-vis *intentions*, this means that we now have a model for loose motivation, namely in terms of the regulative role that we identified by looking at plans in the ordinary sense. Specifically, what we’ve done is to use some form of *strict* motivation to develop a model for *loose* motivation, one that could be used by a proponent of DI to account for the way normative judgement connects to certain responses.

Is this approach feasible? The problem with it is that on closer inspection it is unclear how we should understand it. After all, if an appeal to *plans* is to be helpful in accounting for features of normative judgments, the relevant feature must actually be a feature of plans. Hence, the distinctive regulative role that is being appealed to here must actually be the regulative role of *plans*. However, even if we assume that there are states that play a plan-like role in the regulation of attitudes other than intentions, it is unclear why the way these states regulate other attitudes shouldn’t fall under the model of strict motivation, too. After all, the way *plans* regulate our mental economy *does* fall under that model. Of course, we can *stipulate* that there is a distinctive non-plan like way of structuring, organizing, regulating, and controlling mental states, and that plans—in the ordinary sense—are a special case of states with such a role. But then, it seems, we are merely assuming the features we wanted an account of.

Note that we aren’t arguing that, generally, the strategy of using the basic case of some folk-psychological attitude and abstracting from it to generalize in relevant ways isn’t feasible.¹³ This is a common strategy in philosophical psychology and action theory. A paradigm example is desire. Here philosophers start with the folk-psychological notion of desire and abstract from this familiar state to a more general class of states—which philosophers now tend to label ‘desires’—that share certain crucial features with desires in the ordinary sense.

However, there is a crucial difference between this kind of case and the current proposal for DI. In cases where we legitimately generalize from the particular case to a more general class, what makes the generalization legitimate is that we have good reason to assume that we can fruitfully understand the features possessed by the general class in terms of the relevant features of the individual case. The particular case provides a *model* that allows us to understand the relevant

¹³ Thanks to [omitted] for raising this issue.

features of the more general class. For example, what all desires in the technical sense have in common is that they play a certain role in *intentional explanations*, a feature paradigmatically possessed by desires in the ordinary sense. And so desire in the ordinary case provides an instance of the feature we want to pick out, but also provides us with a model for how to understand the workings of the feature in question. The problem we highlighted for the above account, however, is that these legitimizing features seem absent. What the proponent of DI needs to provide is an informative *model* of loose motivation. And it isn't clear how the suggested account would provide *that*. Let us explain.

When we abstract away from plans in the ordinary sense we aren't starting with a state that loosely motivates, but with a state that strictly motivates. Hence, if our broader sense of plans retains the way plans in the ordinary sense motivate, it seems we can at most generate an account of a general kind of strict motivation, not loose motivation. On the other hand, if we abstract away from the specific way plans motivate, it is difficult to see how we still count as using plans as a model for motivation. It is then no longer the case that we have used the particular case to generalize to the feature we wanted an account of. Hence, such an account wouldn't provide us with a model that would allow us to understand how the internal connection between normative judgements and various responses is supposed to be understood. It wouldn't enhance our understanding of the specific ways such mental states regulate other mental states, or anything along those lines. It would merely tell us that states in the general class regulate mental states somewhat like plans, but not at all in the way plans do.

Hence, it seems that moving to this 'technical' level doesn't advance the dialectic between DMI and the DI: either the features we assume just are the features we find in strict motivation, or we simply stipulate the features we require and thereby fail to provide an account.

6. Conclusion

We started with a choice between two forms of Strict Motivational Internalism, neither of which looked particularly promising. TMI lacked plausibility as an *internalist* thesis, and DMI raised worries about voluntarism. DI looks like a compelling way through the horns of this dilemma. But we've ended by noting that moving from DMI to DI actually reveals a new sort of dilemma.

While DMI faces worries about voluntarism, it comes with resources that can help explain *why* there is a necessary connection between epistemic judgment and motivation. Meanwhile, DI seems to avoid the voluntarism worries, but it does so at an explanatory cost. Unlike DMI, it's not clear that it has the resources needed to explain *why* there is a necessary connection between epistemic judgment and certain doxastic dispositions. Indeed, it seems that the most such an

account can do is restate the view at this point.

The purpose of this paper hasn't been to argue that there is no epistemic analogue of MMI. That would be interesting, because it would seem to tell us something deep and important about the nature of epistemic normativity. It might even suggest that it's a mistake to think of epistemic judgment as genuinely normative after all. As we said above, Hazlett and others have argued *from* claims that epistemic judgment isn't necessarily motivating to the view that epistemic judgment isn't normative. We also pointed out that this would in turn have deep implications for a wide-range of contemporary debates in epistemology, such as debates about the epistemic norm of assertion or epistemic norm of action, or recent work on 'epistemic blame.' If it turns out epistemic judgment isn't normative—or at least not normative in a way that closely parallels moral judgment—we may have reason to re-think such discussions.

But we have examined the landscape and identified what we think is the most promising analogue of MMI, namely DI. We have also identified further theoretical work that needs to be done in order to demonstrate that DI is ultimately viable. Given how much contemporary epistemology seems to rest on an adequate resolution of this issue, this is a worthwhile project.

References

- Adler, J. (2002) *Belief's Own Ethics*. Cambridge, MA: MIT Press.
- Alston, W. (1988) 'The Deontological Conception of Justification', *Philosophical Perspectives*, 2, 257-299.
- Archer, A. (2017) 'Aesthetic Judgements and Motivation', *Inquiry*, 60: 656–674.
- Beddor, B. (2016) *Reduction in Epistemology*, PhD Thesis, Rutgers University.
- Brown, J. forthcoming. 'What is Epistemic Blame? *Noûs*.
- Brown, J. (2008). 'Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning', *Noûs*, 42: 167-189.
- Buckareff, A. (2014) 'Deciding to Believe Redux', in J. Matheson and R. Vitz (eds.) *The Ethics of Belief: Individual and Social*, Oxford: Oxford University Press.
- Church, J. (2002) 'Taking it to Heart: What Choice do We Have?', *The Monist*, 85: 361-380.
- Copp, D. (2001). 'Realist Expressivism: A Neglected Option for Moral Realism', in E. Frankel Paul (ed.) *Moral Knowledge*, 1–44, Cambridge: Cambridge University Press.
- Coté-Bouchard, C. 2017. *Belief's Own Metaethics? A Case Against Epistemic Normativity*, PhD Thesis, Kings College London.
- DeRose, K. (1992) 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research*, 52: 913–929.

- Gibbard, A. (2003) *Thinking How to Live*, Cambridge MA: Harvard University Press.
- Gibbard, A. (2012) *Meaning and Normativity*, Oxford: Oxford University Press.
- Ginet, C. (2001) 'Deciding to Believe', in M. Steup (ed.) *Knowledge, Truth, and Duty*, New York: Oxford University Press.
- Hazlett, A. (2013). *A Luxury of the Understanding*, Oxford: Oxford University Press.
- Hazlett, A. (2014). Expressivism and convention-relativism about epistemic discourse. In A. Fairweather and O. Flanagan (eds.) *Naturalizing Epistemic Virtue*, Cambridge: Cambridge University Press.
- Hieronymi, P. (2009) 'Believing at Will', *Canadian Journal of Philosophy*, Supplementary Volume 35: 149-187.
- Kappel, K., & Moeller, E. F. L. (2014) 'Epistemic Expressivism and the Argument from Motivation', *Synthese*, 191/7: 1529–1547.
- Kauppinen, A. (2018). 'Epistemic Norms and Epistemic Accountability', *Philosophers' Imprint*, 18: 1-16.
- Köhler, S. (2013) 'Do Expressivists Have an Attitude Problem?', *Ethics* 123: 479-507.
- Korsgaard, C. (1996) *The Sources of Normativity*, Cambridge: Cambridge University Press.
- Mitova, V. (2011) 'Epistemic Motivation: Towards a Metaethics of Belief', in A. Reisner & A. Steglich-Petersen (eds.) *Reasons for Belief*, Cambridge: Cambridge University Press.
- Nickel, P. (2010) 'Voluntary Belief on a Reasonable Basis', *Philosophy and Phenomenological Research*, 81: 312-334.
- Owens, D. (2000) *Reason Without Freedom: The Problem of Epistemic Normativity*, London: Routledge.
- Peels, R. (2014) 'Believing at Will is Possible', *Australasian Journal of Philosophy*, 93: 524-541.
- Rettler, L. 2017. 'In Defense of Doxastic Blame', *Synthese*, online first.
- Ridge, M. (2007) 'Expressivism and Epistemology: Epistemology for Ecumenical Expressivists', *Proceedings of the Aristotelian Society*, 81: 83–108.
- Ridge, M. (2014) *Impassioned Belief*, Oxford: Oxford University Press.
- Ridge, M. (2018) 'How to Be an Epistemic Expressivist', In C. McHugh, J. Way, & D. Whiting (eds.) *Metaepistemology*, 141–158, Oxford: Oxford University Press.
- Ryan, S. (2003) 'Doxastic Compatibilism and the Ethics of Belief', *Philosophical Studies*, 114: 47-79.
- Shah, N. (2002) 'Clearing Space for Doxastic Voluntarism', *The Monist*, 85: 436-445.
- Smith, M. (1994) *The Moral Problem*, Oxford: Blackwell.
- Stanley, J. (2005) *Knowledge and Practical Interests*, New York: Oxford University Press.

- Steup, M. (2008) 'Doxastic Freedom', *Synthese*, 161: 375-382.
- Tresan, J. (2006) 'De Dicto Internalist Cognitivism', *Noûs*, 40/1: 143–165.
- Wedgwood, R. (2008) *The Nature of Normativity*, Oxford: Oxford University Press.
- Williams, B. (1973) 'Deciding to Believe', in *Problems of the Self: Philosophical Papers 1956-1972*, Cambridge: Cambridge University Press.