

Evading the Doxastic Puzzle by Deflating Epistemic Normativity*

Luis Oliveira
University of Houston

Introduction

About 30 years ago, William Alston (1988) penned the *locus classicus* for a puzzle that is at the heart of contemporary debates on epistemic normativity. Alston's puzzle, in short, comes from realizing that the most natural way of understanding talk of epistemic justification seems to be in tension with the limited control we have over our belief formation.¹ In this paper, I want to clarify and expand this puzzle, as well as examine the nature and full consequences of a deflationary approach to its resolution.

I begin by abstracting from Alston's discussion in three ways. First, I will shift from his narrow focus on epistemic justification and will focus instead on the broader notion of a *doxastic obligation*: an obligation regarding what one believes, disbelieves, and suspends judgment about.² Alston clearly thinks of the deontological conception of justification as depending on more fundamental facts about our obligations to believe, so the shift in focus I am proposing is not at all a departure from Alston's own way of thinking about this issue. Next, I will shift from a narrow focus on belief and speak instead of our alleged obligations to have various *doxastic attitudes*: belief, disbelief, and suspension.³ This too is in line with Alston's own way of thinking, despite his surface focus on belief. Finally, and still following Alston, I take a cue from the common language typically used to express ideas related to my target notion:

There are many locutions that encourage us to think of believing as subject to requirement, prohibition, and permission. We say "You shouldn't have supposed so readily that he wouldn't come through", "You have no right to assume that", "I had every right to think that she was honest", "I ought to have given him the benefit of the doubt", and "You shouldn't jump to conclusions". (p. 119)

* For comments and discussions, I am grateful to Hilary Kornblith, Chris Meacham, Louise Antony, Peter Graham, Justin Coates, Timothy Perrine, Josh DiPaolo, Ed Ferrier, Dennis Kavlakoglu, Daniel McGloin, Lucas Santos, Heidi Garbe, and the editors of this volume.

¹ Alston is extremely careful in his discussion of a variety of agential notions: basic voluntary control, non-basic immediate voluntary control, long-range voluntary control, and indirect voluntary influence. Nonetheless, since my concern here is not with the voluntarism vs involuntarism debate, I will simply use the term *effective control* as a cover term for whatever key aspect of one's agency is most plausibly a requirement on obligation and most plausibly controversial in the context of belief.

² I will take an *epistemic obligation* to be a specific kind of *doxastic obligation*, one originating from exclusively epistemic sources (the value of truth or knowledge, for example, or the nature of belief, or the nature of the practice of forming beliefs, and so on). Plausibly, one can have non-epistemic doxastic obligations too, originating perhaps from moral or prudential sources. (One can accept this, of course, without endorsing any particular claim about the interaction between these sources and obligations—See Maguire and Woods (forthcoming) for discussion.) At any rate, since my paper discusses a challenge to the existence of doxastic obligations in general, it *ipso facto* discusses a challenge for the existence of epistemic obligations in particular.

³ I am following Friedman (2013) in taking suspension of judgment as a distinctive doxastic attitude

However, I will once again abstract away from these particular claims and focus instead on what I will call general *doxastic ought-claims*: claims whose propositional content is equivalent to “S ought to have doxastic attitude D towards p at t”. My guiding presumption is that the truth of some such claim is behind the truth of the one’s in Alston’s examples. More importantly, my guiding presumption is that the truth of some such claim is behind the truth of most epistemic norms: S ought to believe that p in accordance with S’s evidence; S’s credence in p ought to conform to the axioms of the theory of probability; S ought not to have contradictory beliefs; etc. If the truth of general doxastic ought-claims is problematic, then so is the truth of all these norms.

With these three modifications in mind, I will call what follows *the Doxastic Puzzle*, or *the puzzle* for short:

The Doxastic Puzzle

Each of these claims seems true, but at least one must be false:

Doxastic Deontology (D): Claims of the form ‘S ought to have doxastic attitude D towards p at t’ are sometimes true at t.

Ought-Implies-Can (OIC): If Φ -ing at t is not within S’s effective control at t, then it is false, at t, that ‘S ought to Φ at t’.

Doxastic Involuntarism (I): For all S, p, and t, having doxastic attitude D towards p at t is not within S’s effective control at t.

If (D) is true, then the consequent of (OIC) sometimes is false when Φ -ing stands for *having doxastic attitude D towards p*. And if (I) is true, then the antecedent of (OIC) is always true when Φ -ing stands for *having doxastic attitude D towards p*. Yet (OIC) cannot be true if its antecedent is always true, and its consequent is sometimes false, for the same interpretation of Φ -ing. Alstonian exegesis aside, this broader characterization of the doxastic puzzle is my true concern in this paper.

All three natural replies to the puzzle have been pursued. Some have claimed that doxastic attitudes like believing that p are, in fact, within our effective control, or sufficiently so. Others have claimed that doxastic ought-claims, strictly speaking, are always false. And some have denied that effective control is required for the adequacy of doxastic ought-claims in general.⁴ What all of these different reactions to the puzzle have in common, however, is that they are attempts to *dispel* it: they solve the puzzle by eliminating one of its constitutive claims. Yet given the initial and resilient plausibility of these three claims, none of the various resulting views is theoretically inexpensive: each denial requires either silencing strong intuitions or modifying

⁴ For representative arguments, see Ryan (2003), Steup (2008), and Weatherson (2008) for denials of (I); Sosa (2007), Hazlett (2013), and Papineau (2014) for denials of (D); and Chuard & Southwood (2009), Graham (2011), and McHugh (2012) for denials of (OIC).

central aspects of one's epistemology and/or one's general normative outlook. Each view accepts a claim that is denied by another, after all, and each also denies a claim that is accepted by the other two. I myself am initially inclined to accept all three claims and am resistant to rejecting any of them, despite the extant arguments, upon reflection. So perhaps there is room in this dialectic for a more diplomatic resolution.

In the first part of this paper, then, propose to articulate some such diplomatic resolution. The core suggestion is that (I), (D), and (OIC) are not only each true but actually not in tension with each other in the first place. Instead of attempting to dispel the puzzle, notice, this solution proposes to *evade* it instead: to solve it by properly understanding, and by thereby accepting without contradiction, all of its constitutive claims. I will refer to this as *the evasive strategy* against the puzzle. My discussion examines this evasive strategy in light of recent developments in debates about normativity and the semantics of ought-claims. In the end, however, what is revealed is that such evasion is not a fully diplomatic strategy—it is not without its own surprising consequences. As I argue in the second part of the paper, the evasive strategy forces us to re-think our understanding of the place of normative reasons in epistemology. More exactly, it seems to come at the cost of one central way of thinking about our reasons for having doxastic attitudes, one where such reasons are good-standing exemplars of normative reasons in general. The evasive strategy, that is, threatens to lead us very quickly to what I will call a *deflationary* picture of epistemic normativity: it rescues normative talk, but sacrifices normative substance. I conclude by explaining why I think this is more consequential than some have made it out to be, and by suggesting that these consequences are welcome nonetheless.

1. Evaluations and Prescriptions

Talk of “normativity” is often obscure. Sometimes such talk is about certain *speech acts*; sometimes it is about certain *judgments* expressed or presumed by certain speech acts; sometimes it is about certain *concepts* deployed in certain judgments; sometimes it is about certain *properties* captured by certain concepts. Even worse, it is often unclear what philosophers are indicating when they claim that a speech act, a judgment, a concept, or a property is *normative* in the first place. I want to avoid this kind of obscurity, so I begin by discussing what I mean by talk of normativity in this paper.

However we capture it in our theories, talk of normativity, at bottom, is talk of two familiar kinds of experiences. We experience what I will call *evaluative-normativity* when we evaluate a state of affairs by some standard that we endorse or take for granted. When I say sincerely of a guest that she is “polite”, I am evaluating her behavior with respect to some standard of etiquette that endorse or take for granted (that I have actively or passively accepted). When I say sincerely of an actress that she does not “deserve” to win the award, I am evaluating her acting with respect to some inchoate artistic standard that I endorse or take for

granted as well. In one case, the evaluation is positive; in the other, the evaluation is negative. In both cases, the evaluation is quite familiar.

Ernest Sosa (2007, 70) has in mind precisely these experiences of evaluative-normativity when he says the following:

We humans are zestfully judgmental across the gamut of our experience: in art, literature, science, politics, sports, food, wine, and even coffee; and so on, across many other domains. We love to evaluate even when no practical interest is in play. We judge performances, whether artistic or athletic; grade products of craft or ingenuity; evaluate attitudes, emotions, institutions, and much more.

So when we say that some speech act, judgment, concept, or property is normative, sometimes we simply mean to indicate, however vaguely, its special place within these familiar experiences of evaluative-normativity.

A bit differently, we experience what I will call *prescriptive-normativity* whenever it seems to us that some action or attitude is all-things-considered required of someone or, in short, is simply to be done.⁵ When I say sincerely to a friend that he ought to be faithful to his partner, for example, I don't take myself as merely saying that cheating scores poorly with respect to a standard that I happen to actively accept or passively endorse. I would certainly deny this description if offered to me as an account of what I mean. I do, of course, take myself to be saying that *as well*, but some such claim seems too parochial to capture all of what I am trying to convey. Instead, what is characteristic of our experiences of prescriptive-normativity—and what distinguishes them from our experiences of mere evaluative-normativity—is a perception of certain actions as being “demanded” of us (cf. Street 2012, 44), a perception of “being called on” to behave in a certain way (cf. Thomson 2008, 207), and a sense that some facts simply “settle what to do” (cf. McPherson 2018, 244) in my situation. It is this alleged feature of *morality*, for example, the fact that we experience its claims as prescriptive-normative, that Kant (1788, 86) eulogized (“Duty! Thou sublime and mighty name that dost embrace nothing charming or insinuating but requirest submission”) and that Mackie (1977, 38) took as “utterly different from anything else in the universe” and as ultimately reflecting some of our mistaken beliefs about a purely natural world.⁶

⁵ An *all-things-considered requirement* is the final word about what to do. This is potentially different from a requirement being categorical, or inescapable, or objective. As Foot (1972, 311) has famously argued, many requirements that do not seem to have the “practical oomph” (cf. Joyce 2006) that I am trying to capture can nonetheless be categorical, inescapable, and objective. Setting aside the important task of capturing these experiences descriptively in theory, I here assume the reader has enough of a working sense of what I am talking about when I allude to them.

⁶ The source of the normative authority that we seem to commonly experience, and its very existence in fact beyond just thought and language, is obviously a matter of debate. I don't mean to take a stand on that matter. My only intention, at this point, is to describe and distinguish the common experience that lies behind that debate.

Stephen Grimm (2009, 253–4) is alluding to these experiences of prescriptive-normativity when he says the following:

When we judge a belief to be unjustified or irrational, we seem to be doing more than just evaluating (in this case, in a negative way) the skill or virtuosity of the believer’s performance... To judge someone’s belief to be unjustified or irrational is thus to judge that the person’s attitude towards the content of the belief should be reconsidered, in some apparently binding sense of ‘should.’

So when we say that some speech act, or judgment, or concept, or property is normative, sometimes we mean to indicate something more than merely its special place within our familiar experiences of evaluative-normativity; sometimes we mean to indicate, however vaguely, its special place within our equally familiar experiences of prescriptive-normativity instead.⁷

At bottom, then, talk of normativity (colloquial or academic) is at times talk of our experiences of evaluative-normativity and at times talk of our experiences of prescriptive-normativity. Whether we are advising, criticizing, or deliberating, we sometimes engage in these practices by indicating or noticing our active or passive acceptance of some standard of evaluation, and sometimes we go beyond and indicate or notice something a bit more binding. A theory of normativity, of course, aims at *explaining* what makes it appropriate, if anything, for some speech act, or judgment, or concept, or property to play a role in these familiar experiences and practices, and it proceeds by examining what these items and contexts, and the world in general, would have to be like for any of this to make good sense. But simply *distinguishing* between these two kinds of normativity will suffice for now.⁸

2. The English ‘Ought’

Passages such as the above by Grimm, however, may seem to suggest that the English ‘ought’ is only appropriate in the context of prescriptive-normativity. The passage may suggest, that is, that ‘ought’ always

⁷ Nolfi (2014, 98) and others claim that prescriptions are generally connected to the giving of advice, guidance, or direction. But what I am here calling prescriptive-normativity is more narrowly characterized as the experience or expression of an all-things-considered requirement. Nolfi’s way of talking about prescriptions is in fact orthogonal to the distinction between evaluative-normativity and prescriptive-normativity that I am suggesting since we can certainly give advice, guidance, or direction without implying that the prescribed behavior is required in this way (e.g. “You shouldn’t steal this car; you should steal that car over there instead!”). The same remarks apply to hypological notions about praise and blame: nothing here is presumed about whether they properly belong to our experiences of evaluative-normativity or prescriptive-normativity instead. Finally, Joshua DiPaolo (2019) distinguishes between norms of perfection and norms of compensation, where the former tell us how agents ought to behave if they’re to be perfect and the latter tell us how imperfect agents ought to behave in order to compensate for their imperfections; my distinction is orthogonal to his.

⁸ My labels for this distinction follow Southwood (2016, 17-19), though I do not here endorse his full analysis. See Wodak (2019, 1-5) as well for a helpful discussion of this distinction under a different label: *substantively* normative standards versus *merely formally* normative standards. Other labels for the same distinction include *reason-involving* normativity and *rule-involving* normativity (Parfit 2011, 144-145), *robust* and *formal* normativity (McPherson 2011, 232-233), and *thin* and *robust* normative standards (Côté-Bouchard and Littlejohn, forthcoming). Whatever labels we choose, it is important to recognize, following Chrisman (2008), that evaluative ought-claims sometimes contextually imply prescriptive ought-claims, specially depending on what is taken as shared background in the conversation.

serves to express that something is required or to be done. But this is not the case and is in fact the source of some confusion.

Declarative sentences employing the English ‘ought’—*ought-claims*—can be ambiguous in familiar and uncontroversial ways. An utterance of ‘it ought to rain tomorrow’ sometimes expresses a claim about what would be best to happen (as in ‘it ought to rain tomorrow, if we are to survive this drought’), and it sometimes expresses a claim about what is likely to occur (as in ‘it ought to rain tomorrow, if the readings of the Doppler radar are correct’). According to a widely accepted view, nonetheless, this familiar feature of ought-claims does not reflect underlying semantic differences.

According to this orthodox view, ought-claims always express that some proposition p is true in all the best contextually salient possible worlds, given a contextually salient standard of evaluation.⁹ On Angelika Kratzer’s (1977; 1981) widely used terminology, the contextually salient worlds form a *modal base* (a set of possible worlds consistent with the information in conversational background), and the contextually salient standard of evaluation consists of an *ordering source* (a feature of the conversational background that produces a partial ranking of the worlds in the modal base). Let ‘ W ’ stand for the set of contextually salient possible worlds, and let ‘ S ’ stand for the contextually salient standard of evaluation. We can then say that, according to the orthodox variety of the Uniformity View, the semantic structure of (simple) ought-claims is always “ $O(p)^{W,S}$ ” and it always means something like “ p is true in all the worlds ranked highest in W , given S .”¹⁰

Beside the elegance and plausibility that comes from its simplicity, one of the main advantages of this view is that it allows us to understand the compositional structure of ought-claims, despite its many variations, and therefore explain language acquisition and communication (cf. Chrisman 2012c, 313). Another key advantage, more relevant to our present purposes, is that this view can easily explain the common ambiguities by noticing that different contexts can make salient different kinds of standards of evaluation (epistemic, deontic, prudential, etc.). The same set of contextually salient possible worlds, that is, can be ranked differently, at different times, thereby expressing different meanings.

No doubt there is a use of ‘ought’ closely connected to our experiences of prescriptive-normativity. ‘Ought’, in what we can call *prescriptive-normative ought-claims*, expresses that there is an authoritative relation of normative requirement or obligation holding between a certain agent S and a certain action Φ . This is likely what we mean when we tell a tempted friend that he ought to be faithful to his partner. Prescriptive-normative ought-claims, in other words, function by ranking salient possible worlds according to some

⁹ For some proponents, see Phillips (1998), Wedgwood (2007), Chrisman (2012a, 2012c), and Finlay & Snedegar (2014).

¹⁰ I am simplifying things a bit since the variations within this view are irrelevant for my arguments in this paper.

authoritative standard of evaluation—a standard that is guided by values that “call for certain responses and make us criticisable if we’re not invested or committed to serving” them (cf. Côté-Bouchard and Littlejohn, forthcoming, 6-7). But there are uses of the English ‘ought’ closely connected to our experiences of evaluative-normativity as well. ‘Ought’, in what we can call *evaluative-normative ought-claims*, merely expresses the ranking of a state of affairs with respect to a non-authoritative standard that is being accepted or taken for granted. This is likely what we mean when we tell a visiting friend that she ought to take the second exit on the left: we do not thereby express a normative requirement, we rather express what ranks best according to the familiar standards of efficient traveling—standards that we reasonably presume to be shared and salient in this context. The English ‘ought’ is thus appropriate in contexts of prescriptive-normativity as well as in contexts of evaluative-normativity, depending on what it is used, at that time, to express: a relation of normative requirement between agents and actions, or *merely* a ranking-relation between standards and states of affairs.¹¹ Expressing the latter—it is worth emphasizing the key point—is not always the same as expressing the former.

This feature of the English ‘ought’ is quite important. Since evaluative-normative ought-claims and prescriptive-normative ought-claims express different kinds of relations, they can be expected to have different truth conditions. In particular, it is plausible that only evaluative-normative ought-claims can be true when the grammatical subject of the sentence is not capable of bringing about the relevant state of affairs (cf. Chrisman 2008, 56-7). Consider the claims:

(A) The clock ought to strike on the quarter hour.

(B) Wealthy Americans ought to donate to Oxfam.

Since clocks have no agency, and since striking on the quarter hour is not an action which clocks can choose to perform or not, (A) would be false or incoherent if it claimed that there is a relation of normative requirement holding between a certain agent (clocks) and a certain action (striking on the quarter hour). Yet (B) seems coherent and even true when taken in that way. Wealthy Americans are agents, after all, and we typically think of them as perfectly capable of choosing between alternative actions. This “agency condition” to the truth of some but not all ought-claims, we can now say, depends on the kind of standard that is being employed by the ought-claim. The *surface-grammar* similarity between (A) and (B), in other words, does not reflect a *normative similarity*. While (A) is simply an evaluative-normative claim, (B) is also a prescriptive-normative claim; while (A) is simply an evaluation of a certain state of affairs, (B) is also a claim about a

¹¹ Southwood (2016) distinguishes between four different ‘oughts’. My distinction is broader and encompasses his.

certain relation of normative requirement that holds between certain agents and a certain action. As a consequence, claims like (A) can be true in kinds of situations where claims like (B) cannot.¹²

The English ‘ought’ is thus an instructive example of the importance of being clear about our talk of normativity. I have suggested that, at bottom, such talk is about our familiar experiences of what I have called evaluative-normativity and prescriptive-normativity. What the English ‘ought’ illustrates, however, is that what makes a speech act, a judgment, a concept, or a property appropriate in the context of one of these experiences may well differ from what makes it appropriate in the context of the other. Depending on which standard of evaluation is employed, different kinds of claims, with different truth-conditions, are expressed.

3. The Evasive Strategy

For our present purposes, we can turn the doxastic puzzle into an argument against (D):

1. If Φ -ing at t is not within S 's effective control at t , then it is false, at t , that ‘ S ought to Φ at t ’.
2. For all S , p , and t , having doxastic attitude D toward p at t is not within S 's effective control at t .
- C. So claims of the form ‘ S ought to have doxastic attitude D toward p at t ’ are never true at t .

This argument seems to show that a commitment to (OIC) and (I) forces a rejection of (D). But this appraisal is premature. We have already noted the surface-grammar ambiguity between evaluative-normative ought-claims and prescriptive-normative ought-claims. Now we must clarify which of these two kinds of ought-claims are being deployed in each of the premises. In general, very many debates in normative philosophy would be clearer if this fact were foregrounded and more energy were consequently focused on determining which kind of normativity is at stake.

The first stage of the evasive strategy is precisely the suggestion that true claims of the form ‘ S ought to have doxastic attitude D toward p at t ’ are always evaluative-normative ought-claims and never prescriptive-normative ought-claims. When we say truly that S ought to have doxastic attitude D toward p , that is, we are always expressing something of a kind with what we express by saying that the clock ought to strike on the quarter hour; we are always merely claiming that S 's having doxastic attitude D toward p at t is well-ranked by a certain *epistemic* standard that we accept or take for granted, and we are never claiming that there is a relation of normative requirement holding between S and the having of D at t . As Côté-Bouchard and Littlejohn (forthcoming) put it:

¹² See Southwood (2016, 23-28) for the claim that “prescriptive oughts” (in his sense) do not always require effective control. His argument, however, depends both on a particularly narrow sense of “control,” as well as on a classification of prescriptive oughts as something belonging to the category I am calling evaluative-normativity.

Unlike moral facts, epistemic facts are only thinly normative. Unlike moral norms, and just like etiquette, fashion, and the law, epistemic norms are such that there is not necessarily a robust reason to conform to them.

In precisely this sense, Kornblith (2001, 238) takes true doxastic ought-claims as expressions of our *epistemic ideals*, Chrisman (2008, 364; 2012b, 603) takes them as *rules of criticism* or *state-norms*; and Feldman (2000, 676; 2008, 351) takes them as *role oughts*. Though different in their details, these are all accounts of why doxastic ought-claims are not prescriptive-normative but evaluative-normative instead.

For clarity, let ‘ought_E’ stand for the evaluative-normative sense of the English ‘ought’ and let ‘ought_P’ stand for the prescriptive-normative sense. We can now say that, according to the evasive strategy, the proper understanding of (D) is the first of these and not the second:

Evaluative-Normative Doxastic Deontology (D_E): Some claims of the form ‘S ought_E to have doxastic attitude D toward p at t’ are true.

Prescriptive-Normative Doxastic Deontology (D_P): Some claims of the form ‘S ought_P to have doxastic attitude D toward p at t’ are true.

If this interpretation is correct, and if we accept the suggestion in the previous section, then the truth of claims of the form ‘S ought to have doxastic attitude D toward p at t’ does not require that S be capable of bringing about the state of affairs of having D. D_E can be true, that is, even if S lacks effective control over having D at t.

The second stage of the evasive strategy, then, is the suggestion that (OIC) cannot be true if its consequent refers to evaluative-normative ought-claims. Having effective control, after all, is simply not required for their truth. According to the evasive strategy, that is, the proper understanding of (OIC) is the first of these and not the second:

(OIC_P): If Φ -ing at t is not within S’s effective control at t, then it is false, at t, that ‘S ought_P to Φ at t’.

(OIC_E): If Φ -ing at t is not within S’s effective control at t, then it is false, at t, that ‘S ought_E to Φ at t’.

According to Southwood (2016, 10), in fact, the ambiguity behind careless statements of (OIC) explains our ambivalence towards its truth:

Whether or not “ought” implies “feasible” depends on which core practice-serving ought is at issue. There is some salient core practice-serving ought for which it is true that “ought” implies “feasible.” There is some other salient core practice-serving ought for which it is false that “ought” implies “feasible.” The interpretation of those normative claims that seems false on account of demanding the infeasible—and for which it seems objectionably unworldly to insist otherwise—involves the former. The interpretation of the claims that seems true in spite of demanding the infeasible—and for which it seems unduly lenient to insist otherwise—involves the latter. Our ambivalence, far from involving any mistake, is entirely appropriate.

Properly understood, therefore, (OIC) is the claim that a certain relation of normative requirement holds between an agent and an action only if that action is within that agent's effective control.¹³ Properly understood, (OIC) places no constraints on the truth of any positive or negative mere evaluation. This is what Feldman (2008, 347-8) has called *modest deontologism*.

The third and final stage of the evasive strategy simply replaces the original premises of the relevant argument with their appropriate disambiguations:

1*. If Φ -ing at t is not within S 's effective control at t , then it is false, at t , that ‘ S ought _{t} to Φ at t' ’.

2. For all S , p , and t , having doxastic attitude D toward p at t is not within S 's effective control at t .

C*. So claims of the form ‘ S ought _{t} to have doxastic attitude D toward p at t' ’ are never true at t .

(C*), however, is not the denial of (D_E). While (I) and (OIC_P) show that no prescriptive-normative ought-claims are true about an individual and her doxastic attitudes, they nonetheless spell no trouble for (D_E).

Whatever we say about S 's agency, sometimes it is true that S believing that p , for example, is well ranked according to the typical epistemic standards that we actively or passively accept or take for granted.

The evasive strategy thus amounts to accepting the conjunction of the following claims:

The Evasive Strategy: $(I) \wedge (OIC_P) \wedge (D_E) \wedge \sim(OIC_E) \wedge \sim(D_P)$

This is why evasion, to me, seems to be the least controversial and the least costly of the available maneuvers against the doxastic puzzle: it does not require a rejection of (I), it does not require a wholesale rejection of (OIC), and it does not require abandoning normative talk about what we ought to believe.

¹³ Thus Côté-Bouchard and Littlejohn (forthcoming, 8) “Animals are not subject to moral standards. Your dog’s actions cannot be morally right or wrong. More generally, since animals lack the rational capacities needed for agency and attributions of responsibility, they are exempt from robustly normative standards. They cannot have robust reasons that apply to them and the “oughts” that require such reasons do not apply to them.”

4. Evasion and Normative Reasons

Despite appearing to be a more diplomatic resolution to the puzzle than its alternatives, and despite its full or partial widespread favor, the evasive strategy is not without its own substantive implications. On one widely accepted picture of the relation between normative reasons and normative requirements, the evasive strategy sacrifices the very existence of normative reasons for being in a doxastic state. I will outline the four-step argument for this claim in this section.

The first step is accepting the familiar picture of normative reasons as the building blocks of relations of normative requirement: it is precisely *because* S has undefeated normative reasons in favor of Φ -ing that an authoritative relation of normative requirement holds between S and Φ -ing. While our reasons are considerations of various strengths in favor of various and sometimes conflicting actions at a certain time, the balance of these considerations determines what we are required to do. Using standard terminology, the idea here is that the prescriptive-normative English ‘ought’—‘oughtp’, the ‘ought’ that expresses a relation of normative requirement between an agent and an action—indicates what we have *all-things-considered* reasons, or *decisive* reasons, to do. Consider Parfit (2011, 33):

When we have decisive reasons, or most reason, to act in some way, this act is what we *should* or *ought* to do in what we can call the *decisive-reason-implying* senses. Even if we never use the phrases ‘decisive reason’ or ‘most reason’, most of us often use ‘should’ and ‘ought’ in these reason-implying senses.

Parfit’s view of the connection between normative reasons and what I have been calling prescriptive-normativity is echoed by McPherson (2011, 233): “almost all of us take reasons to be normative in a way that contrasts with other formally normative systems such as the rules of chess.” While the rules of chess ground various evaluations that we at times might take quite seriously, they do not, by themselves, provide us with the kind of normative reason that can ground an authoritative requirement.

We can call the view behind these claims *normative rationalism*:

Normative Rationalism (NR): There is a normative requirement for S to Φ if and only if (and because) there is most normative reason for S to Φ .¹⁴

Accepting (NR), notice, does not require taking a stand on the controversial question of whether considerations other than evidence can count as a reason to have doxastic attitude D. It also does not require

¹⁴ This is one way of expressing the view made famous by W.D. Ross (1930) and further developed by Nagel (1970), Scanlon (1998), Raz (2002), and, as quoted, Parfit (2011). Notice that (NR) could be stated as a weaker, *satisficing* principle by substituting ‘most normative reason’ for ‘sufficient normative reason’. My argument in this section goes through on either formulation. For some resistance to the import of this meta-normative view to the epistemic domain, see Sylvan and Sosa (2018).

taking a stand on whether these reasons are reasons to promote some value or to respect some norm. (NR) is simply an account of the source of normative requirements, of the kind of normativity that we experience as binding. As such, (NR) is initially compatible with whatever views one has about what counts as reasons for doxastic attitudes, as well as whatever views one has about the sources of these reasons. It is important not to confuse one's reasons for rejecting one of these other views with a reason to reject (NR). All that said, I think (NR) is true and will take it on board in what follows. Those who are unsure or positively resistant can restate my argument in conditional form.

The second step is noticing an implication of (NR): if there is no relation of normative requirement holding between S and Φ -ing, then either (a) there are no normative reasons for S to Φ , or (b) these reasons are *defeated* by reasons in favor of doing something else that is incompatible with Φ -ing, or (c) there are normative reasons for S to Φ , reasons that are undefeated by reasons for doing something else that is incompatible with Φ -ing, but reasons that nonetheless are not strong enough to ground a relation of normative requirement between S and Φ -ing. At nearly all times, there are available actions such that it is not the case that S ought_t to perform them at that time. Sometimes this is the case because S has no reason whatsoever to perform them. Perhaps right now, for example, I have absolutely no reason to raise my index finger straight up into the air; consequently, it is not the case that I ought_t to raise my finger in that way: raising my finger, right now, is not something that I am required to do. But sometimes it is not the case that S ought_t to perform a certain action despite in fact having reasons for performing it. Perhaps right now, for example, I have a reason to have a cup of coffee (because I have a non-conflicted desire for it), but have even stronger reasons not to have it (because I've already had too many in a short amount of time). So while I have a reason to have a cup of coffee, it is not the case that I ought_t to have a cup of coffee: having a cup of coffee is not something that I am required to do. Finally, sometimes it is not the case that S ought_t to perform a certain action despite in fact having undefeated reasons for performing it. Perhaps I have reasons to go help my neighbor set-up his garage sale, reasons that are not defeated by my reasons for doing anything else right now, but reasons that are nonetheless too weak to ground a *requirement* to go across the street. It's just not true that I ought_t to go help him with his chairs. The point here, however, is that one of these three explanations is true whenever we have the absence of a relation of normative requirement between S and Φ -ing.

The third step is noticing an implication of the conjunction of (NR) and the evasive strategy's denial that claims of the form 'S ought_t to have doxastic attitude D toward p at t' are *ever* true. Given the implication of (NR) just mentioned as my second step, this denial means accepting that it is *always* true that either:

- (a1) There are no normative reasons for S to have doxastic attitude D toward p at t.

(b₁) S's reasons for having doxastic attitude D toward p at t are defeated by S's reasons in favor of doing something else that is incompatible with having doxastic attitude D toward p at t.

(c₁) S's undefeated reasons for having doxastic attitude D toward p at t are not strong enough to ground a requirement to believe that p.

Since reasons are the building blocks of relations of normative requirement, denying that a certain relation of normative requirement holds between S and Φ means denying that there are *decisive reasons* for S to Φ .

Denying that there is *ever* a relation of normative requirement holding between S and Φ , therefore, means denying that there are *ever* decisive reasons for S to Φ (because of (a) or (b) or (c) considerations). But this calls out for explanation.

The fourth and final step, then, is realizing that consideration (a) is always the best explanation for why there is no normative requirement for S to have doxastic attitude D toward p at t. It would certainly be odd if every case where S has a normative reason for having doxastic attitude D toward p were also a case where that reason is defeated by S's reasons for doing something else that is incompatible having D toward p. The only thing that seems incompatible with believing that p, for example, is yet another doxastic attitude: not believing that p or suspending judgment about p. So the only way to have reasons for having some doxastic attitude D that are defeated by decisive reasons that are incompatible with having D is to have decisive reasons for another doxastic attitude and, consequently, a normative requirement to have it. Similarly, it would certainly be odd if every case where S has an undefeated normative reason for having doxastic attitude D toward p were also a case where those reasons are not strong enough to be decisive. If there are no decisive reasons for believing that p, once again, and no decisive reasons for believing that not-p, then *ipso facto* there are decisive reasons for suspending judgment about p.¹⁵ So there are doxastic normative requirements for having some doxastic attitude D toward p in that case as well, no matter what. So the best explanation for the inexistence of any normative requirement for having doxastic attitude D toward p, and perhaps even the only explanation, is that no one ever has normative reasons for having doxastic attitude D toward p at all. There are simply no such things.

We can state my argument in this section more precisely in the following way:

The No Reasons Argument

¹⁵ This may seem to be in tension with *permissivism*: the view that “some bodies of evidence permit more than one rational doxastic attitude toward a particular proposition” (Jackson and Turnbull, forthcoming). But this is not the case. Permissivism, strictly speaking, is a view about “evidence”, and not a view about “normative reasons.” Any sense of tension here depends on additional substantive commitments. I am grateful to Justin Coates for bringing this to my attention.

1. If the evasive strategy is correct, then claims of the form ‘S ought_p to have doxastic attitude D toward p at t’ are always false.
2. If claims of the form ‘S ought_p to have doxastic attitude D toward p at t’ are always false, then there are never decisive normative reasons for having doxastic attitude D toward p.
3. If there are never decisive normative reasons for having doxastic attitude D toward p, then there are no normative reasons for having doxastic attitude D toward p.
- C. So if the evasive strategy is correct, then there are no normative reasons for having doxastic attitude D toward p.

Premise 1 follows from the fact that the evasive strategy accepts both (I) and (OIC_p). This is an important element of the evasive strategy, since it thereby avoids some of the costs of alternative strategies against the doxastic puzzle. Premise 2 follows from the familiar picture of normative reasons as the building blocks of requirement-relations: (NR). Premise (3) follows from the traditional three-way exclusive division of doxastic attitudes and the structural connections between any possible reasons for or against them.

It is here that we see the most important implication of the evasive strategy: while it succeeds at escaping the doxastic puzzle by taking doxastic oughts as evaluative-normative ought-claims that do not require agency, it sacrifices the existence of normative reasons for having doxastic attitudes. While it retains one kind of normativity for belief, it abandons the kind of normativity that we experience as binding. As I will put it, it thereby *deflates* epistemic normativity by removing any of its prescriptive bite. Evading the doxastic puzzle by deflating epistemic normativity, as we now can see, is not as fully diplomatic as it looked on first inspection.¹⁶

Conclusion

I have argued that the evasive strategy, despite the widespread appeal of its constitutive parts, comes at the cost of normative reasons for having doxastic attitudes. The significance of this consequence cannot be overstated. Since talk of reasons is ubiquitous in epistemology, the evasive strategy forces sympathizers to find something ameliorative to say about such talk. Perhaps, for example, talk of “epistemic reasons” is talk of a merely domain-relative kind of reason (cf. Sosa 2007, Ch. 4) and, as such, not talk of full-blown, authoritative, normative reasons after all. Perhaps, for another example, talk of “having evidence for p” is not the same as (and it doesn’t imply) talk of “having a reason to believe that p” (cf. Leite 2007), such that we can retain talk of evidence while abandoning talk of reasons. Both maneuvers, I should say, seem good to me.¹⁷

¹⁶ Here is another potential cost of the evasive strategy: Perhaps S is responsible for believing that p only if there are normative reasons for S to believe that p. There is nothing implausible about this claim. As Skorupski (2010, 295) puts it: “there can be no reason to blame someone who had no reason not to do what he did.” But if (NR) and this claim are both true, then the evasive strategy sacrifices not only the existence of normative reasons for belief, but also the existence of *epistemic responsibility*. I won’t pursue that line of argument here, but see Perrine (forthcoming) for a recent discussion of the doxastic puzzle and matters of responsibility.

¹⁷ I have argued elsewhere against versions of evidentialism stated in terms of prescriptive ought-claims (Oliveira 2017, 2018a), and against internalist notions of justification stated in agential terms (Oliveira 2015, 2018b).

Yet it is important to note one central way in which they fail to be fully ameliorative, relative to the full implications of the evasive strategy.

Consider my fellow deflationary epistemologists Côté-Bouchard and Littlejohn (forthcoming). They argue that “our epistemic speech can correctly pick out epistemic facts even if there are no robustly normative facts and reasons” (p. 5), and that “epistemic norms... are thinly normative, not robustly so” (p. 7). This is very much in the spirit of the evasive strategy and the deflated epistemic normativity I have outlined above.

However, they go on to argue that nothing of importance is really lost by deflating epistemic normativity in this way. This is because we can still do and say everything that normativity and normative talk is supposed to allow: we can still ground epistemic evaluations, these evaluations can still be important, we can still talk in terms of epistemic permissions, epistemic standards can still be inescapable, and we can still talk in terms of epistemic reasons. What more could we want from epistemic normativity? And yet while this is very much in the spirit of my view, there is a deeper loss to deflating epistemic normativity in the ways I have suggested.

Deflating epistemic normativity in the wake of evasion, as I will put it, rules out the possibility of *normative upgrade*. And normative upgrade, in turn, is the key ameliorative feature that makes deflation palatable to most others. We see this, for example, when Côté-Bouchard and Littlejohn compare epistemic normativity to the normativity of etiquette: “the connection between belief and epistemic standards seems more akin to the connection between arrangements of cutlery and the standards of etiquette than the connection between actions and moral standards” (p. 9). The suggestion, of course, is that, like the reasons of etiquette, epistemic reasons, despite being themselves only thinly normative, can also be upgraded: “We think that the force of the criticism that we level against a thinker for failing to believe what she ought to believe comes from non-epistemic reasons—reasons to form beliefs that meet epistemic standards” (p. 9). Just like you can have a true binding requirement to follow the thinly normative rules of etiquette on some extraordinary occasion (whenever authoritative normative reasons enjoin you to so do it) the suggestion is that you can likewise have a true binding requirement to follow the thinly epistemic norms as well. Consider another pair of deflationary epistemologists, Maguire and Woods (forthcoming, 6), on this very same point:

We often have practical reasons to have correct beliefs. And here, epistemic reasons play a key role. The interaction of epistemic reasons explains which belief is correct in a given situation. This is analogous to other activities. Chess-based reasons – that this move will win material, or that that move will cost you a tempo, are not authoritative either. They don’t explain what you just plain ought (henceforth: ought) to do. Practical reasons do that. But we sometimes have practical reasons to make correct chess moves – where by this we mean: make the strategically best move within the rules of the game. The interaction of chess-based reasons explains which move is correct. This structure preserves a key normative role for

epistemic reasons for beliefs, and a key normative role for practical reasons for beliefs. Each plays a distinctive role in the explanation of what we ought to believe.

With the possibility of normative upgrade, that is, indeed very little is lost by deflating epistemic normativity. But this is precisely where the deflationary epistemology produced by the evasive strategy reveals its cost. As the No Reasons argument makes clear, the evasive strategy does not merely dampen the authority of a certain *source* of normative reasons for having doxastic attitudes, but rather it eliminates the very existence of normative reasons for having doxastic attitudes in the first place. It does this, more exactly, by relying on the “agency condition” for the truth of certain ought-claims. On the resulting picture, then, epistemic normativity is not at all like the normativity of chess and etiquette, both of which can be upgraded when other kinds of reason are at play, but rather like the normativity of clocks (as suggested by Chrisman) or digestion (as suggested by Alston). The evasive strategy, as we can now see, deflates epistemic normativity all the way down to its lowest or thinnest level, and it is incompatible with the kind of ameliorative normative upgrade that other deflationary epistemologists take for granted.

With that in mind, I should say that I myself am attracted to a reconstructive reply to the consequences of evasion: I take it as an invitation to excise the notion of an *epistemic reason* from its center place in epistemology. Kornblith (2015), for example, argues that talk of “reasons” is a relic from epistemological theorizing that is overly focused on the phenomenology of deliberation, anyway, and is lacking attention to the psychology of belief fixation. I think something along those lines is promising. Externalist theories such as process reliabilism, for example, can provide us with the tools to understand the central epistemic notions (knowledge and justification) without any reference to epistemic reasons, and we have just seen that we can recover talk of “ought to believe” without them just as much. Of course, there is much more to be said about this. Nonetheless, I think there is good reason to see what would happen to an epistemology that truly put talk of “epistemic reasons,” and corresponding talk of “epistemic obligations,” to the side.

References

- Alston, William (1988). “The Deontological Conception of Epistemic Justification,” *Philosophical Perspectives* 2: 257-299.
- Chrisman, Matthew (2008). “Ought to Believe,” *Journal of Philosophy* 105 (7): 346-370.
- Chrisman, Matthew (2012a). “Ought and Control,” *Australasian Journal of Philosophy*, 90(3): 433-451.
- Chrisman, Matthew (2012b). “The Normative Evaluation of Belief and the Aspectual Classification of Belief and Knowledge Attributions,” *Journal of Philosophy* 109(10): 588-612.
- Chrisman, Matthew (2012c). “On the Meaning of ‘Ought’” In Russ Shafer Landau, ed., *Oxford Studies in Metaethics*, Vol 7: 304-333. Oxford University Press.

- Chuard, Philippe & Southwood, Nicholas (2009) "Epistemic norms without voluntary control," *Noûs* 43(4): 599-632.
- Côté-Bouchard, Charles, and Littlejohn, Clayton (forthcoming). "Knowledge, Reasons, and Errors About Error Theory," in *Metaepistemology: Realism and Anti-Realism* by Christos Kyriacou & Robin McKenna (eds.). Palgrave Macmillan.
- DiPaolo, Joshua (2019). "Second Best Epistemology: Fallibility and Normativity," *Philosophical Studies* 176 (8): 2043-2066.
- Feldman, Richard (2000). "The Ethics of Belief," *Philosophy and Phenomenological Research* 60(3): 667-695.
- Feldman, Richard (2008). "Modest Deontologism in Epistemology," *Synthese*, 161, 339-355.
- Finlay, Stephen & Snedegar, Justin (2014). "One Ought Too Many," *Philosophy and Phenomenological Research*, 86(1): 102-124.
- Foot, Philippa (1972). "Morality as a System of Hypothetical Imperatives," *Philosophical Review*, 81 (3): 305-316.
- Friedman, Jane (2013). "Suspended Judgment," *Philosophical Studies*, 162 (2): 165-81.
- Graham, Peter A. (2011). "Ought and Ability," *Philosophical Review*, 120(3): 337-382.
- Grimm, Stephen (2009). "Epistemic Normativity," in *Epistemic Value*, by Haddock, Miller, & Pritchard (eds.): 243-264. Oxford University Press.
- Hazlett, Allan (2013). *A Luxury of the Understanding: On the Value of True Belief*. Oxford University Press.
- Hieronymi, Pamela (2008). "Responsibility for Believing," *Synthese* 161(3): 357-373.
- Jackson, Elizabeth, and Turnbull, Margaret Greta (forthcoming). "Permissivism, Underdetermination, and Evidence," in *Routledge Handbook for the Philosophy of Evidence*, by Clayton Littlejohn & Maria Lasonen-Aarnio (eds.): 1-13. Routledge.
- Joyce, Richard (2006). *The Evolution of Morality*. MIT Press.
- Kant, Immanuel (1788/1956). *Critique of Practical Reason*, trans. Lewis White Beck. University of Chicago Press.
- Kornblith, Hilary (2001). "Epistemic Obligation and the Possibility of Internalism," in *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, by Fairweather & Zagzebski (eds.): 43-264. Oxford University Press.
- Kornblith, Hilary (2015). "The Role of Reasons in Epistemology." *Episteme* 12: 225-239.
- Korsgaard, Christine (1996). *The Sources of Normativity*. Cambridge University Press.
- Kratzer, Angelica (1977). "What Must and Can Must and Can Mean," *Linguistics and Philosophy* 1: 337-55.
- Kratzer, Angelica (1981). "The Notional Category of Modality," in Eikmeyer, H. and Rieser, H. (eds), *Semantics From Different Points of View*: 117-47. Berlin: Springer.
- Leite, Adam (2007). "Epistemic instrumentalism and reasons for belief: A reply to Tom Kelly's 'epistemic rationality as instrumental rationality: A critique'," *Philosophy and Phenomenological Research* 75 (2): 456-464.
- Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. Penguin.
- Maguire, Barry, and Woods, Jack (forthcoming). "The Game of Belief," *Philosophical Review*.
- McHugh, Conor (2012). "Epistemic Deontology and Voluntariness," *Erkenntnis* 77(1): 65-94.
- McPherson, Tristram (2011). Against Quietist Normative Realism. *Philosophical Studies*, 154(2), 223-240.

- McPherson, Tristram (2018). "Authoritatively Normative Concepts," in *Oxford Studies in Metaethics, Volume 13* by Russ Shafer-Landau (ed.): 253-278.
- Nagel, Thomas. (1970). *The Possibility of Altruism*. Princeton University Press.
- Nolfi, Kate (2014). "Why is Epistemic Evaluation Prescriptive?" *Inquiry* 57(1): 97-121.
- Oliveira, Luis R.G. (2015). "Non-Agential Permissibility in Epistemology," *Australasian Journal of Philosophy* 93 (2): 389-394.
- Oliveira, Luis R.G. (2017). "Deontological Evidentialism, Wide-Scope, and Privileged Values," *Philosophical Studies* 174 (2): 485-506.
- Oliveira, Luis R.G. (2018a). "Deontological Evidentialism and Ought Implies Can," *Philosophical Studies* 175 (10): 2567-2582.
- Oliveira, Luis R.G. (2018b). "Ampliative Transmission and Deontological Internalism," *Pacific Philosophical Quarterly*. 99 (2): 174-185.
- Papineau, David (2014). "There Are No Norms of Belief," in *The Aim of Belief* by Chan (ed.): 64-80. Oxford University Press.
- Parfit, Derek (2011). *On What Matters*. Oxford University Press.
- Perrine, Timothy (forthcoming). "Strong Internalism, Doxastic Involuntarism, and The Costs of Compatibilism," *Synthese*.
- Phillips, David (1998). "The Middle Ground in Moral Semantics," *American Philosophical Quarterly* 35(2): 141-155.
- Raz, Joseph (2002). *Engaging Reason*. Oxford University Press.
- Ross, W. D. (1930/2002). *The Right and the Good*. Clarendon Press
- Ryan, Sharon (2003). "Doxastic Compatibilism and the Ethics of Belief." *Philosophical Studies* 114 (1-2): 47-79.
- Scanlon, T.M. (1998). *What We Owe Each Other*. Harvard University Press.
- Skorupski, John (2010). *The Domain of Reasons*. Oxford University Press.
- Sosa, Ernest (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford University Press.
- Southwood, Nicholas (2016). "Does 'Ought' Imply 'Feasible'?" *Philosophy and Public Affairs*, 44 (1): 7-45.
- Steup, Mathias (2008). "Doxastic Freedom," *Synthese* 161(3): 375-392.
- Street, Sharon (2012). "Coming to Terms with Contingency: Humean Constructivism About Practical Reason," in *Constructivism in Practical Philosophy* by Lenman & Shemmer (eds.). Oxford University Press.
- Sylvan, Kurt, and Sosa, Ernest (2018). "The Place of Reasons in Epistemology," in *The Oxford Handbook of Reasons and Normativity*, by Daniel Starr (ed.). Oxford University Press: 555-575.
- Thomson, Judith J. (2008). *Normativity*. Open Court.
- Weatherston, Brian (2008). "Deontology and Descartes' Demon," *Journal of Philosophy* 105 (9): 540-569.
- Wedgwood, Ralph (2007). *The Nature of Normativity*. Oxford: Oxford University Press.
- Wodak, Daniel (2019). "Mere Formalities: Normative Fictions and Normative Authority," *Canadian Journal of Philosophy*, 49: 828-850.

Forthcoming in *Epistemic Duties: New Arguments, New Angles* (Routledge),
by Kevin McCain and Scott Stapleford (eds.)