

Metaethical Deflationism, Access Worries and Motivationally Grasped Oughts

Sharon Berry

Abstract

Mathematical knowledge and moral knowledge (or normative knowledge more generally) can seem intuitively puzzling in similar ways. For example, taking apparent human knowledge of either domain at face value can seem to require accepting that we benefited from some massive and mysterious coincidence. In the mathematical case, a pluralist partial response to access worries has been widely popular. In this paper, I will develop and address a worry, suggested by some works in the recent literature like [3], that connections between ought facts and action prevent us from giving a similarly pluralist response to moral access worries.

1 Introduction

Mathematical knowledge and moral knowledge (or normative knowledge more generally) can both seem intuitively puzzling in similar ways. In both cases, we take ourselves to have many true beliefs. But, taking this claim at face value can seem to require positing a spooky coincidence. Given that we cannot see, touch, or otherwise causally interact with mathematical objects, why should there be any relationship between what mathematical principles people accept and what abstract mathematical objects actually exist? And how could it be anything but a mysterious coincidence that our moral beliefs line up with objective moral facts (as traditionally understood)?

In the case of mathematics, it is popular to give a partially *metasemantic* pluralist[17] response to access worries, along the following lines. Say that (for one reason or another) almost any logically coherent pure mathematical sentences we could have adopted as axioms would have come to express truths, so that (given accuracy about logical coherence) our acceptance of truth-preserving pure mathematical axioms no longer seems mysterious. And then avoid claiming our mathematical practices are especially fitting in some other way that could revive the appearance of commitment to a mysterious coincidence[7, 20, 22, 8, 1, 3]¹ In this way we can answer mathematical access

¹For example, neo-Carnapian versions of this metasemantic pluralist approach maintain that if we'd adopted

worries – or at least reduce them to corresponding access worries about our knowledge of logical coherence.

Could we treat moral and mathematical knowledge analogously? One might hope to similarly answer or reduce moral access worries by

- Maintaining that if we had applied “ought” in a different logically coherent way (while retaining current connections between acceptance of “ought” claims and action alongside certain other quasi-analytic principles) we’d have still qualified as speaking the truth about some suitable other notion ought*.
- Rejecting all claims that our moral practices are specially fitting² in any way that would revive intuitive access worries (as a philosophical stance[6]). So, for example, don’t claim that our moral practice is uniquely correspondent with facts about supernatural posthumous rewards and punishment.

Such a deflationary metasemantic answer to moral access worries is clearly incompatible with many forms of traditional moral realism – whose intuitive appeal centers on embracing exactly the kind of special truth-conduciveness or fittingness claims rejected in the two bullet points above [1, 11, 3, 4]. However, it might be useful to metaethical deflationists (like Humean Sentimentalists), who use access worries to motivate their favored deflationary (but still truth value realist) view of moral properties³. For such metaethical deflationists need to make it plausible that switching from traditional moral realism to their view helps answer/avoid access worries. And accepting the two points above suggests a way of doing so - by maintaining that if I’d used moral terms differently but coherently I would have expressed a truth about a correspondingly different moral concept (e.g. about permissibility’ as opposed to permissibility)⁴

different logically coherent pure mathematical axioms, we’d have meant something suitably different by expressions like ‘there is’ and the relevant mathematical symbols, so these alternative mathematical axioms would still have expressed truths.

²c.f. [16]

³Here I have in mind deflationary views, on which people with different sensibilities and dispositions to make practical choices (e.g., normal humans and psychopaths or paperclip maximizing robots) will tend to would express correspondingly different concepts with their moral-talk-like practices.

⁴C.f. the fact that if I’d used words differently in certain ways (e.g., applying the word ‘even’ to numbers of the form $2k + 1$ and odd to numbers of the form $2k$), I would have meant something different by ‘even’ but still expressed

However, recent works by Justin Clarke-Doane[3] and Vermaire[23] raise a worry that answering moral and mathematical access worries analogously in this way prevents us from acknowledging certain intuitive connections between ought beliefs and action. In this paper, I will develop and attempt to answer this worry by suggesting that ‘all things considered ought’ is (what I’ll call) a *motivationally grasped concept*. Doing this lets us give a metaethical deflationist answer to moral access worries without having to reject intuitions about close connections between ought judgments and actions (contra Clarke-Doane and Vermaire worry) or accept controversial claims about logically coherent sociopaths (as some other Hume-inspired metaethical deflationist views do).

In §2 I’ll present the main worry to be considered: a claim that metaethical deflationists (who answer access worries as above) can’t accept or adequately account for the intuitive internalist connections between accepting ‘I ought to ϕ ’ and (something like) intending to ϕ ⁵. Clearly, one possible response would be to simply reject these internalist principles. However, I’ll argue that if we don’t want to do that, we can explain the truth of these internalist principles in a way that’s compatible with the form of metaethical deflationism sketched above. In §3, I’ll propose a theory of ‘all things considered ought’ as (what I’ll call) a *motivationally grasped concept*, which is intended to provide such an explanation.

In §4 I’ll suggest a way of extending this proposal about all things considered oughts to include an answer to moral access worries, via a theory of moral oughts as indirectly motivationally-grasped notions. And in §5 I’ll answer worries that accepting the latter proposal prevents one from making adequate sense of current practices of debate and disagreement over ought.

2 Worries about Ought Belief-Action Connections

In [3] Clarke-Doane raises a worry for the metaethical deflationism sketched above (alongside various other truthvalue realist views) as follows.

a truth. Intuitively considering this fact can help dispel access worries that my knowledge that even numbers are divisible by two and odd numbers are not, requires a mysterious coincidence.

⁵Note that I won’t try to argue for the existence of any important connections between ought talk and action in this paper. Rather, I will allow for the sake of argument that some such connection exists, and try to answer worries (of the sort raised by Clarke-Doane) that adopting the metasemantic approach to moral access worries advocated above prevents one from acknowledging it (or making sense of it).

If moral (or all things considered) ought facts exist, these facts would have to be able to tell us what to do in a way that would end deliberation. But, he says, considering the conceivability of variant ‘oughtlike concepts’ ought* and ought** which pair a given agent with different responses to a situation (somehow) shows ought facts can’t thus end deliberation.

[E]ven if we all things considered ought to kill the one to save the five [in the trolley problem], we all things considered ought* not (for some all things considered ought like notion, ought*). And now the practical question arises whether to do what we all things considered ought, or all things considered ought*, to do.

Various interpretations can be given to this rather brief remark. However, the version of the worry which seems clearest and most troubling to me, goes as follows.

Intuitively, it would be (at least) odd for someone to accept that they ought to ϕ , while (unconflictedly) lacking all intention to ϕ . For example, we generally expect that someone who accepts ‘I all things considered ought to drive to Texas’ will intend to drive to Texas (absent further factors like weakness of the will or perhaps participation in a tragic moral dilemma ⁶). So many find a principle along the following lines attractive⁷.

Ought Judgement Internalist Principle: It would be somehow odd for me to judge that I ought to ϕ ⁸ while not intending[†] to ϕ .

Accepting such an Ought Judgment Internalist (OJI) principle raises a prima facie challenge for the metaethical deflationist view suggested in the introduction to this paper. For, one might fear that such deflationism prevents one from adequately explaining the connection between ought beliefs and intentions[†] it asserts.

⁶Thanks to REDACTED for drawing my attention to the latter possibility

⁷Here, in order to avoid contentiously analyzing just what disjunction of states (e.g., intending to ϕ or experiencing akresia or taking oneself to be facing a tragic moral dilemma or...) is thus expected to accompany believing one all-things-considered ought to ϕ , I’ll just grant (for the sake of argument) that some such expectations are common – and I’ll use the term ‘intending[†] to ϕ ’ to stand for whatever intending-like state is expected to accompany acceptance that one-all-things considered ought to ϕ .

⁸More specifically, it would be odd for a subject to judge that they ought to ϕ without having a corresponding intention[†] under the mode of presentation they’d express by saying, ‘I ought to ϕ ’. There’s nothing odd about a person judging that they ought to ϕ under the description ‘the person who will next meet the Russian spy ought to ϕ ’ while unconflictedly lacking anything like an intention to ϕ .’

For example, a traditional moral/practical normative realist might explain the OJI’s claimed connection between ought beliefs and intentions[†] by saying that recognizing ought facts is *intrinsically motivating for all agents*, regardless of their desires⁹. This motivating power could explain why people who believe they ought to ϕ tend to either intend to ϕ or feel some conflict about their lack of intention to ϕ .

However, advocates of my metasemantic solution to moral access worries plausibly cannot say the same. For, suppose that we say that all logically coherent variants on my ought attribution practices will express truths (about a correspondingly different concept). Then it may be that I have one ought-like concept, ought, and Immanuel has another, ought*, such that I ought to switch the path of a trolley in a certain situation, but I ought* not to switch the path of the trolley in that same situation. If I say that *only my notion (ought)* has intrinsic motivating power, then it looks like I’m taking my ought concept to be special in a way that would intuitively revive access worries. In principle, I could say that both ought and ought* facts have the relevant motivating power (so that knowing that I ought to switch the trolley but ought* not to would require having conflicting intentions[†]). But such a view is *prima facie* unappealing^{10 11}.

Thus, one might fear that accepting and explaining the Ought Judgment Internalist principle above prevents one from adopting the form of metaethical deflationism advocated in the introduction.

⁹Or at least they can say this if they can adequately answer classic ‘queer facts’ [18] worries about whether any beliefs could be intrinsically motivating in this way.

¹⁰Compare this point to Vermaire’s point in [23]

¹¹One might argue (thanks to REDACTED for suggesting a version of this point) that the possibility of tragic moral dilemmas suggests that a linguistically competent person can simultaneously recognize ought-like facts that impose incompatible demands on action; doing so just requires a certain internal tension, characteristic of facing a tragic moral dilemma. This would, technically, block the argument above by showing it’s possible to simultaneously hold incompatible intentions[†]. However, saying this doesn’t really help with the problem above. For, given any situation, there are presumably a range of different logically coherent “ought” application practices analogous to our own (and hence different oughtlike concepts) that require all possible responses to that situation. But, we presumably don’t want to allow that recognizing the truth about how all these concepts apply would reveal every situation to warrant the same kind of internal tension as a tragic moral dilemma.

3 Motivationally Grasped Oughts

3.1 Basic Proposal

In the rest of this paper, I will attempt to answer the above worry. I'll propose that we can explain the Ought Judgement Internalist principle, as follows.

Our grasp on (all-things-considered) ought-like notions partly consists in our disposition to obey certain linguistic rules – including one that connects ought beliefs to intentions. This rule says (roughly) ‘don’t accept that you ought to ϕ when you don’t intend[†] to ϕ ’¹². And we can obey this rule by deciding what to do and then adjusting our ought talk to fit our intentions[†] – rather than vice versa. In this way, one might say that our all-things-considered ought concepts are **motivationally grasped**¹³.

To explain by analogy, consider our understanding of the first order logical connectives. Plausibly, we grasp such logical concepts partly by being disposed to obey certain inference rules. For example, one might say we grasp the concept of conjunction, partly in virtue of accepting something like the following wide scope rule¹⁴.

\wedge rule: ‘Don’t simultaneously accept both A and B and reject $A \wedge B$ ’.

Similarly, I propose that we grasp an (all-things-considered) ought-like concept partly by accepting a certain collection of linguistic rules including (but perhaps not limited to) the following¹⁵.

¹²In other words, the Ought Judgment Internalist principle holds because there are quasi-analytic/conceptually central rules connecting ought talk to intentions, not because of some kind of special motivating power of grasping ought facts.

¹³But note that appeal to intending[†] makes obedience to the Ought Judgement Internalist principle compatible with saying things like ‘I ought to ϕ , but won’t because I’m too cowardly’

¹⁴As has often been noted, it’s not plausible that a narrow scope version of this norm (requiring everyone who accepts both A and B to form/maintain the belief that ‘A and B’) has this status. For, grasping the meaning of ‘and’ doesn’t require you to sit around and attempt to close your beliefs under conjunction, by drawing inference after inference to form new beliefs. It doesn’t even require you to infer that ‘A and B’ whenever you believe both A and B and the question of whether ‘A and B’ is raised. For ceasing to believe A or B is also an option.

¹⁵Because this paper is intended as a mere advertisement for a research program, I won’t commit myself to specific view about which other principles and rules (like perhaps the principles of deontic logic covered in [19]) are necessary or conceptually central to our grasp of ‘all things considered ought’ concepts.

Relatedly, I won’t propose a theory about exactly which variants on our actual “all things considered ought” application practices would count as employing variant ought-like concepts, as opposed to using the word “ought” to mean something entirely different. Because my aim is to solve/reduce access worries that our ought accuracy requires some mysterious coincidence, I just need to show certain variants on our actual practices would be equally truth conducive (and refrain from taking our practices to be special in some other ways). Questions about which of these variant *truth conducive* practices count as employing a variant *oughtlike concept* matter little.

(Approximate) Ought Rule: Don't accept 'I ought to ϕ ' while not intending[†] to ϕ .

Importantly, this Ought Rule has a wide-scope form, analogous to that of the ' \wedge ' rule above. That is, it doesn't require anyone who currently believes they ought to ϕ to form a corresponding intention[†] to ϕ . Rather, it merely requires us to (somehow) avoid the state of simultaneously believing we ought to ϕ while not intending[†] to ϕ .

Accordingly, we can explain the Ought Judgment Internalist principle above by appeal to the conceptual centrality of this Ought Rule, rather than by claiming ought beliefs have a mysterious universally motivating force. We can say that the extension of each person's 'ought' concept reflects what 'ought' claims they're disposed to (stably) accept – and hence (given their disposition to obey the Ought Rule) what they're disposed to (stably) intend[†]. In particular, we might think of the situation as follows.

- My long-run dispositions to apply 'all things considered (ATC) ought' reflect a combination of first-person ought claims (fitting what I'm disposed to intend[†] given suitable information and reflection) with certain conceptually central platitudes (claims I treat as would-be analyticities), that might include axioms for deontic logic or claims like 'One ought not sanction a person for taking an action, unless that person (ATC) ought to avoid that action.'
- Almost any ¹⁶ logically coherent such pattern of dispositions to apply 'ought' (given sufficient time for deliberation) will determine a corresponding extension (set of \langle agent, action \rangle pairs, across all metaphysically possible worlds), for my word 'ought' to (rigidly) apply to¹⁷, with the following features.

– I can reliably¹⁸ form true beliefs about what I ought to do (i.e., reliably deploy the

¹⁶See the point about conservativity above.

¹⁷There may be some indeterminacy in the application of the resulting concept, in cases where we aren't disposed to make a judgment either way.

¹⁸I propose that changing your ought beliefs to match your intentions[†] is just a reliable way of forming true beliefs (not an infallible one), for the following reason. The extension of your all things considered ought concept will reflect how you're disposed to apply the term given full descriptive knowledge of a situation and plenty of time for reflection where you might change a view you initially adopted – just as the extension of your term 'bachelor' reflects how you're disposed to apply the term given sufficient time for reflection and retraction. Thus first person all things considered ought facts will tend to reflect what we are disposed to intend[†] given sufficient time for reflection. Accordingly, when we adjust ought beliefs to match hastily-formed intentions[†] (i.e., ones we'd be disposed to change, given more descriptive knowledge or time for reflection), we can wind up with false beliefs.

concept expressed by first-person applications of *my* word “ought” by **forming an intention[†] about what to do** (that’s suitably robust under reflection and knowledge of descriptive facts) and then adjusting the ought claims I accept accordingly.

– I can reliably form true beliefs about what third parties ATC ought to do (in my sense of the word “ought”, by forming beliefs that follow from (or at least harmonize with) first-person oughts and various platitudes which I treat as would-be analyticities (more on this below).

- People who (logically coherently) accept sufficiently different ought platitudes (i.e. are disposed to treat different ought sentences as quasi-analytic) or have sufficiently different intentions[†] will use the phrase ‘all things considered ought’, to express a concept with a correspondingly different extension.

Admittedly, deciding what to do and updating your all-things-considered ought beliefs to match might seem like an odd way of learning ought facts. However, I think it’s not so different from something we already acknowledge: the special access a person has to their own beliefs via what some philosophers call transparency reasoning – the process involved in, e.g., going from accepting ‘the sky is blue’ to accepting ‘I think that the sky is blue’. Just as you can reliably learn whether you believe the sky is blue by asking yourself whether the sky is blue (even though ‘the sky is blue’ does not entail ‘I believe that the sky is blue’), I’m suggesting you can reliably learn what you all-things-considered ought to do by settling (i.e., forming a suitably descriptively informed and reflectively stable intention[†] about) what to do, and then updating your “ought” beliefs accordingly.

3.2 Contrast With Other Views

This proposal differs from current work claiming acceptance of an ought claim has both a belief and an intention/desire component or my acceptance of ‘I ought to ϕ ’ (in some particular case) is *partly constituted by* having the corresponding intention[†]. For I want to allow that it’s possible (though rare) for people to believe they ought to do something without having any corresponding

intentions^{†19}. Above I’ve suggested that having a motivationally-grasped all things considered ought concept requires being *disposed* to obey the Ought rule above (connecting ought acceptance to intentions[†]). But having this disposition to harmonize ought beliefs with intentions[†] doesn’t imply successful performance in each particular case, just as having the disposition to add correctly doesn’t imply perfect performance in any specific instance. Indeed, attention to such occasions of failure to harmonize intentions[†] with ought beliefs (i.e., cases where they serenely accept an ought claim without having the corresponding intention or any internal struggle etc.) might be important for understanding the history of inaction on climate change, global injustice etc.²⁰²¹

My proposal also differs from claims that facts about agents’ desires or choice dispositions *ground* ought facts. I want to reject commitment to any such grounding claim. There’s good independent reasons to think we don’t need to accept any such grounding claims to give a (partially) *metasemantic response* to access worries about knowledge of ought facts — which expects a close relationship between my actual world ‘ought’ application dispositions and the extension of my ‘all things considered ought’ concept (i.e., the set of $\langle agent, action \rangle$ pairs across metaphysically possible worlds). For example, presumably, if I’d used words differently in certain ways (e.g., applying the word ‘even’ to numbers of the form $2k + 1$ and odd to numbers of the form $2k$) I would have meant something different by ‘even’ but still expressed a truth. And considering this fact can help dispel access worries (about my knowledge that even numbers are divisible by two and odd numbers are not). But, clearly, these divisibility facts aren’t **grounded in** or made true by anything about

¹⁹That is, it’s possible to believe you out to do something without having an intention to do that thing or an internal conflict as per weakness of the will or a tragic moral dilemma etc.)

²⁰Compare also Kierkegaard’s remark that, “In relation to their systems most systematizers are like one who builds an immense castle and lives in a shack nearby: they do not live in their own gigantic systematic buildings.”[13]). We might study what prompts these failures to align our ought beliefs with action just as we study what prompts arithmetic failures.

²¹In [9] Horgan and Timmons propose a cognitivist expressivism, which differs from classic expressivism by maintaining that ought beliefs are a form of genuine belief. On this view, having ought beliefs is a *sui generis*, irreducible psychological state of ought-commitment to a state of affairs which contrasts with *is*-commitment (and there are no “metaphysically robust moral facts”). The state of ought commitment is (alongside other phenomenological features), “inherently motivational; one is typically motivated to judge and act accordingly (although [this] motivational strength can be outweighed”. Although I find much to sympathize with in Horgan and Timmons’ view, I disagree with them by favoring a looser connection between ought beliefs and motivation. On their view, it should be impossible for ought beliefs to come apart from motivation. In contrast, as noted above, I find it somewhat more attractive to say that ought beliefs without corresponding motivation are not impossible but just rare and pathological like (erroneous) refusals to accept instances of $\&I$ or $\&E$ (because having a motivationally grasped ought concepts requires something like a disposition to obey the Ought Rule, but this disposition need not always be realized).

human language use (c.f., Boghossian on epistemic vs. metaphysical analyticity[2]).

This proposal differs from simple response dependence theorists[12], in not claiming that any specific principle of connecting meaning to use dispositions (“‘I ought to ϕ ’ is true iff I’d be disposed to accept ‘I ought to ϕ ’ in such-and-such descriptively characterized situation”) is an analytic truth, obvious to all competent English speakers. I am committed to the idea that meaning partly reflects (something like) ideal use dispositions. But the latter claim doesn’t imply that any specific metase-mantic theory precise enough to enable reductive definition (‘x ought to ϕ iff I would judge that he ought to ϕ under such-and-such-conditions’) is obvious to all linguistically competent speakers.

Finally this proposal importantly differs from one of its closest rivals in the literature (Street’s [21] Humean constructivism) in avoiding controversial internalist commitments regarding third person claims.

In [21] Sharon Street proposes that facts about what an agent has reason to do (and hence, presumably, what they all-things-considered ought to do) reflect something like what would best satisfy their ideally coherent desires. She also maintains facts about an agent’s reasons are correspondingly grounded in facts about their desires. So, for example, Street controversially holds that an ideally coherent Caligula all-things-considered ought to murder people. This claim is unpopular (or at least very controversial). My proposal might seem to commit one to this controversial view.

Similarly, imagine a deeply apathetic person, Pierre, who constantly, non-conflictedly, feels no motivation to act²². My proposal might seem to implausibly imply that, e.g., it’s not the case that Pierre ought to save a drowning child.

My proposal (as stated so far) avoids this problem, by implying very few constraints on what we can third personally claim Pierre ought to do. We are free to answer access worries by saying that almost²³) any logically coherent way of combining my first-person ought attribution dispositions (including dispositions to obey the Ought Rule above) with quasi-analytic principles connecting first and third-person ought claims would wind up being truth conducive (by securing a suitable meaning

²²So, for example, he doesn’t even feel wistful regret that he lacks strong first-order values, or anxiety that he’s sinning against some highly valued but poorly grasped standard he doesn’t know how to learn more about

²³We might also require that the resultant ‘ought’ practice be conservative over much of the rest of your language[5], in the sense of, e.g., not letting one prove any new claims statable using only non-normative vocabulary (or any such claims that didn’t already express truths). But this requirement tends to be easily satisfied by practices that are logically coherent.

for ‘ought’). Thus my story is compatible with saying that people with common dispositions to intend[†] to praise and sanction can truly say (using their motivationally grasped all things considered ought concept) “Pierre ought to save a drowning child (despite the fact that saving a child is nothing to him)”.

Admittedly, my proposal *does* suggest that a (sufficiently) deeply apathetic Pierre couldn’t *himself* truly say “I ought to save that drowning child”, while using ought to express motivationally grasped ought concept. However, accepting this isn’t much of a bullet to bite, for (as I suggest in the example about a sociologist and then expand on below) there are various important and common non-motivationally-grasped ‘ought’ usage practices Pierre could fully participate in (associated with advice crowdsourcing, persuasion and historical/psychological explanation). So lingering impulses to say that surely apathetic Pierre could still grasp the fact that ‘he ought to save the drowning child’, might be understood by reference to these.

However, some readers may have a further worry that my story leaves it uncomfortably mysterious *why we would ever care* to make and hear claims about what third parties like Caligula ought to do. We often spend time thinking and talking about what third parties all things considered ought to have done (e.g., when gossiping, judging and backseat driving). Yet the story I’ve told so far has little to say about such third person all things considered ought claims. So one might fear that it implausibly suggests such third person talk will tend to be meaningless, confused, or unimportant.

I’ll address this worry as part of a story about moral (as opposed to all things considered) oughts in the next section.

4 Moral Oughts

In this section I’ll propose a rough but (hopefully) motivating example of how the story about access to all things considered ought facts advocated above could be extended to an answer to access worries about moral oughts.

Speaking a bit abstractly, I want to suggest we accept certain core platitudes which function

somewhat like a ramified implicit definition, in Frank Jackson's proposal in [10]²⁴ One might our grasp of moral oughts and our grasp of theoretical terms in psychology and cognitive science (like the superego or innate grammar modules). In both cases, one might say, we grasp theoretical terms by (something like) associating them with certain core theoretical posits/would-be platitudes²⁵, which pin down the intended role these psychological/moral properties are supposed to play in a larger theory that predicts and explains the application of antecedently understood terms.

In the psychological case, these platitudes might specify how (say) superego facts are supposed to help predict and explain more concrete psychological and behavioral facts. But (I want to suggest) the moral case, these platitudes would express how moral oughts are supposed to combine with other antecedently understood facts to predict (and perhaps, in a sense, explain) motivationally grasped first person ATC oughts.

For example, the platitudes associated with a typical English speaker's moral ought concept might include the following.

- If something is morally obligatory, then I all-things-considered ought to do it.
- If A isn't morally obligated to ϕ then I morally ought (and hence all-things-considered ought) to not blame A for failing to ϕ (or punish them for it or demand compensation for it).

Supposing that we accept some conceptual principles like this can explain why my all things considered ought claims about what a third party like Caligula all things considered ought to do are significant and interesting. For (while accepting the above expected platitudes, maintaining logical coherence and obeying the Ought rule above) I can't deny that Caligula ought to eschew murder while simultaneously intending[†] to blame him or impose certain punishments on him for murdering²⁶.

²⁴However (to use Quine's metaphor), the periphery – which the internal parts of this web of beliefs aims to match – doesn't just include observation statements but also motivationally grasped first person all things considered oughts.

²⁵Here I use the term platitudes as a reference to Frank Jackson's [10], and the idea that conditional claims like 'If there is a superego then it does something like ϕ ' might feel trivial/conceptually central/quasi analytic in cases where ϕ is one of these theoretical platitudes.

²⁶If (while accepting the platitudes above) I say that Caligula all-things-considered can murder, I must also be willing to accept that Caligula morally can murder. And (according to the same platitudes above) one morally (and hence all-things-considered) ought not to blame/punish someone for doing things they morally can do. So if I say that Caligula all-things-considered can murder (and I accept all the platitudes above), then I must accept that I

On the picture I'm sketching, one can generally reliably form true beliefs about what third parties ought to do by combining knowledge of what *we* all-things-considered ought to do (got via forming intentions and the Ought rule above) with platitudes connecting moral oughts to all-things-considered oughts, like the ones above.

Note, however, that I don't claim ought platitudes are unassailable and must express truths. For (depending on what you care about and what intentions[†] you're disposed to have) it's not guaranteed that some extension can be found for your term 'morally ought' (and other such moral vocabulary) which satisfies all relevant platitudes and lets these terms play their intended theoretical role in predicting all-things-considered-ought facts. Compare this to how conceptually central principles about the superego or an innate grammar module might fail to express truths if there's no choice of extensions for these terms which lets them play their intended/conceptually core theoretical role in predicting and explaining your more observable psychological behavior.

Relevant moral vocabulary terms will only non-defectively apply if there's some possible choice of extension for the theoretical term which makes enough of the relevant platitudes come out true (without unduly changing the meaning/extension of antecedently grasped terms). For example, there's probably no way of assigning extensions to practical-normative concepts like noblesse oblige, the white man's burden, manliness etc. which both satisfies relevant platitudes involving these terms and lets them fulfill their intended theoretical role of helping predict/track the kinds of facts expressed by first person 'all things considered ought' facts in the mouth of a typical modern person. So these terms (if directly used by us in ordinary contexts²⁷) generally won't non-defectively apply.

I take it to be a contingent psychological fact (perhaps explained by things like Kitcher's proposal in [15]) that most of us are disposed to practically deliberate in ways that allow 'morally ought' to be assigned an extension that makes common platitudes come out true (without requiring significant change to non-moral vocabulary).

all-things-considered ought not to punish him. And I can only do the latter (compatible with the Ought Internalist principle above) if I intend[†] to not punish Caligula for murdering.

²⁷I don't mean to rule out a successful sociological/deferential use of these terms (analogous to sociological use of 'all things considered ought' mentioned above) which is divorced from any expected connection to motivationally grasped all things considered ought facts.

5 Disagreement

I will end this paper by considering a pair of worries about the intelligibility of our current moral and all things considered ought disagreement practices.

5.1 Presumption of Shared Subject Matter?

First, one might worry that my proposal makes common responses to apparent disagreement about moral or all things considered ought claims like “Caesar ought not to have crossed the Rubicon” look irrational. Does my proposal imply the implausible conclusion that (as Street[21] puts it) “it is impossible for you and me sensibly to disagree about whether X is a reason to Y for A, since the answer might be ‘yes’ for me but ‘no’ for you”?

My main response to this worry is to deny that two parties need mean exactly the same thing by ‘ought’ to have a meaningful disagreement (or rationally respond to apparent disagreement over oughts in the way we ordinarily do). Someone who accepts my proposal can explain the rationality of conducting moral arguments *as if* all parties expressed the same concept with their ‘ought’ talk, as follows. In normal contexts it is rational to (defeasibly) assume that all parties practically deliberate in *sufficiently similar* ways for their moral and all things considered ought concepts to be nearly coextensive – so that the particular sentence “X ought to A” (and, to some extent, other sentences which come up in relevant arguments) is *very likely* to have the same truth value in all disputants’ ideolects.

Absolute confidence that a given utterance expresses exactly the same thing in my mouth and yours isn’t needed to justify (largely) familiar practices of mutual questioning and argument over ought claims. Nor is it needed to justify beliefs that disputants genuinely do have *some* disagreement. For note that I don’t need to know whether your concept of ‘restaurant’ applies to some weird edge cases like food trucks in the same way mine does, to reasonably have an argument about “whether the building housing our old favorite record store has now turned into a shoe store or a restaurant”. I only need to have sufficiently high justified confidence that the business in

question either counts as a restaurant on both of our favored senses of the term or neither²⁸.²⁹

Additionally, my metaethical deflationism can borrow familiar ideas from expressivists for explaining our moral disagreement practices. For example, the expressivist can note that people seemingly making incompatible moral claims tend to have a *practical disagreement*/disagreement in action³⁰, even if there's no proposition whose truth value they disagree on. Also philosophically minded disputants who accept some form of metaethical realism, can have a genuine further disagreement whenever they seem to disagree about 'whether agent A ought to ϕ ', as follows. They agree in (falsely) believing there's a unique notion which both of their 'all-things-considered ought' talk refers to, and then disagree about how this unique favored notion applies.

5.2 Moral vs. Mathematical Disagreement

A second worry for my proposed metaethical deflationism concerns differences between common responses to apparent disagreement in mathematical vs. moral discourse. In the introduction to this paper, I suggested it was attractive to treat moral and mathematical access worries similarly, given the various similarities between the two cases which have been noted in the companions in innocence literature. However, one might note that there's also an important disanalogy. We (currently) seem much more willing to accept logically coherent variants of our mathematical practices as expressing truths than logically coherent variants on our mathematical practices. Thus, (arguably) current mathematical practice much more obviously fits with and suggests metasemantic pluralism than current moral practice does. Does recognizing this difference significantly reduce the motivation for my project of answering moral and mathematical access worries analogously (via appeal to metasemantic pluralism)?

²⁸Indeed, arguably we should never be completely confident that our interlocutors mean exactly the same proposition by some sentence under dispute as we do, given e.g. the existence of regional variants of English.

²⁹To further motivate this point, consider a thought experiment inspired by ring species in biology. Imagine a chain of 500 cities speaking very similarly, but with a gradual shift in the application of the term "fuchsia" as one travels north, such that "fuschia" is applied to red things in the first city and blue things in the 500th city. If the members of each city only talk to people in neighboring cities, their color talk might proceed just as smoothly as ours does. Since neighbors could be reasonably presumed to apply "fuchsia" similarly, acting on the defeasible presumption that something is fuchsia in my neighbor's sense if and only if it's fuchsia in my sense is a reasonable policy. So we can explain the fact that familiar ways of conducting arguments about whether something is "fuchsia" are rational and fruitful without taking both sides of these debates to be using the term to pick out exactly the same extension.

³⁰By this, I mean that both parties are trying to achieve incompatible aims.

I want to acknowledge the relevance of disanalogy, but suggest two points that I think soften the blow.

First, we may be able to explain the differences between common moral and mathematical disagreement behavior in a way that's compatible with my proposal as follows. I've already suggested that the core motivationally-grasped ought talk discussed in this paper exists alongside other important kinds of ought talk (employed in, e.g., advice-giving, advertising³¹, persuasion and sociology). When using 'ought' in the latter ways we really do mean to pick out the same extension as our interlocutor. So we may be able to explain some of the differences in common responses to moral vs. mathematical disagreement (and corresponding different immediate appeal of pluralist/metasemantic answers to access to worries in domains) by noting the existence of such multiple uses of ought, which are not sharply distinguished in current practice³².

Second, it's worth noting that some of our current tolerant/pluralist attitude to disagreement in mathematical practice (which fuels metasemantic pluralist responses to access worries) appears to be a relatively recent mathematical advance (cf. Kitcher's story about resistance to accepting imaginary numbers in [14]). So perhaps some of this difference in attitude to apparent disagreement in both domains would disappear if the metaethical deflationist project I'm proposing in this paper succeeds!

³¹In the advertising case, we try to motivate someone to act in a certain way by highlighting descriptive facts, which we hope will relate to their methods of practical deliberation to produce a certain action. We try to get people to consider certain descriptive facts in the context of deciding specific practical questions. So, for example, we might say something like '...so that's why you ought to take some vacation, and here's why that vacation ought to be a cruise with Royal Caribbean'. Using "ought" deferentially to express something coextensive with our audience's ought concept is helpful for doing this. In the case of advice-giving, someone deliberating about whether to do something invites input (mention of descriptive facts that might help them end deliberation) from people who know their motivational structure fairly well.

Regarding advice-giving, I would also suggest that there can be interesting ambiguity about whether a person is looking for help deliberating simpliciter, or also expecting answers that reflect facts about your third person all things considered ought judgments and plans (e.g., seeking information about what they can do without risking certain kinds of sanctions from you). For example, this kind of ambiguity might come into play when an omnivore asks a vegan for advice about planning a frugal shopping list, or when a head of government asks an economist or communications expert for advice. The advice-asker may be asking about what they ought to do relative to their own values (more carefully: inviting facts and proposals to be mentioned that will combine with their own methods of deliberation to settle the question of what to do). However, they may instead be asking a question seeking information about what attitudes the *advisor* takes or intends to take (involving praise, blame, sanction etc.) e.g. seeking information about what the advise-seeker can do without risking certain kinds of sanction.

³²Perhaps it is unsurprising that these heterogeneous uses are often not distinguished, given the points about our practical deliberations and motivationally grasped ought concepts being similar to most interlocutors, for it to be a good working hypothesis that extensional differences won't matter to any particular question under consideration.

Thus, overall, I think the theory of motivationally grasped oughts advocated above has significant resources to explain our current practices of handling apparent moral (and all-things-considered) ought disagreement, even in cases where this might initially seem quite challenging.

6 Conclusion

In this paper, I've defended a metaethically deflationist answer to moral access worries, which treats moral and mathematical access worries analogously in some important regards. I've tried to show (via a theory of motivationally grasped ought concepts) that this approach avoids certain bad consequences. In particular, accepting this approach doesn't prevent you from accepting and explaining connections between ought judgments and actions (as recent works by recent works by Justin Clarke-Doane[3] and Vermaire[23] might suggest) or require you to make controversial claims about what a logically coherent Caligula ought to do (like Street's Humean constructivism[21]).

References

- [1] Sharon Berry. (Probably) Not companions in guilt. *Philosophical Studies*, 175(9):2285–2308, 2018.
- [2] Paul Artin Boghossian. Analyticity Reconsidered. *Nous*, 30(3):360–391, 1996.
- [3] Justin Clarke-Doane. *Morality and Mathematics*. Oxford University Press, 2020.
- [4] Matti Eklund. *Choosing Normative Concepts*. Oxford University Press, 2017.
- [5] Hartry Field. *Science Without Numbers: A Defense of Nominalism*. Princeton University Press, 1980.
- [6] Bas C. van Fraassen. *The Empirical Stance*. Yale University Press, October 2008.
- [7] Geoffrey Hellman. *Mathematics Without Numbers*. Oxford University Press, USA, 1994.

- [8] Eli Hirsch. *Quantifier Variance and Realism: Essays in Metaontology*. Oxford University Press, 2011.
- [9] Terry Horgan and Mark Timmons. Grippled by Authority. *Canadian Journal of Philosophy*, 48(3-4):313–336, 2018.
- [10] Frank Jackson. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford University Press, 1998.
- [11] Silvia Jonas. Mathematical and Moral Disagreement. *Philosophical Quarterly*, 70(279):302–327, 2020.
- [12] Antti Kauppinen. Moral Sentimentalism. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2022 edition, 2022.
- [13] Søren Kierkegaard. *Fear and Trembling*. New York: Cambridge University Press, 1986.
- [14] Philip Kitcher. *The Nature of Mathematical Knowledge*. Oxford University Press, 1983.
- [15] Philip Kitcher. *The Ethical Project*. Harvard University Press, March 2014. Google-Books-ID: SWyW0XB7pN8C.
- [16] Stephanie Leary. Choosing normative properties: a reply to Eklund’s Choosing Normative Concepts. *Inquiry*, 63(5):455–474, 2020.
- [17] Øystein Linnebo. Epistemological Challenges to Mathematical Platonism. *Philosophical Studies*, 129(3):545–574, 2006.
- [18] J. L. Mackie. *Ethics: Inventing Right and Wrong*. Penguin Books, 1977.
- [19] Paul McNamara and Frederik Van De Putte. Deontic Logic. In Edward N. Zalta and Uri Nodelman, editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2022 edition, 2022.

- [20] Stuart Shapiro. *Philosophy of Mathematics: Structure and Ontology*. Oxford University Press, USA, 1997.
- [21] Sharon Street. Constructivism About Reasons. *Oxford Studies in Metaethics*, 3:207–45, 2008.
- [22] Amie L. Thomasson. *Ontology Made Easy*. Oup Usa, 2015.
- [23] Matthew Vermaire. Against Schmought. *The Journal of Philosophy*, 118(9):465–485, September 2021.