

WHAT ABILITY CAN DO

ABSTRACT

One natural way to argue for the existence of some subjective constraint on agents' obligations is to maintain that without that particular constraint, agents will sometimes be obligated to do that which they lack the ability to do. In this paper, I argue that while promising, this strategy, unsupplemented, is doomed to fail. Specifically, I argue that because the truth of an ability ascription depends on a (usually implicit) characterization of the relevant possibility space, different metaethical accounts take obligation to be constrained by different *senses* of ability. As a result, what initially looks to be a point of consensus—that ability constrains obligation—turns out to be a point of contention, and arguments with this at the foundation are much more likely to obscure, rather than resolve, metaethical disputes. Despite this, I insist, appeals to ability in metaethics might still bear fruit. If we can independently establish a particular sense of ability as normatively relevant, then we have good grounds for ruling out metaethical accounts that are inconsistent with it. In the final section, I make a preliminary attempt to do just that, arguing that the correct (i.e., normatively relevant) sense of ability is both epistemically and motivationally restricted.

INTRODUCTION

Here's a natural thought: what we're obligated to do is constrained by what we're able to do. In other words,

Ought-Implies-Can (OIC): If an agent ought to ϕ , then she is able to ϕ .

This idea is both intuitive and widely endorsed by philosophers of otherwise disparate metaethical bents. As such, it appears to be a promising starting point for a metaethical debate—if abilities constrain obligations, then any general constraints on what we're able to do will also constrain what we're obligated to do. And this, in turn, suggests a particularly attractive strategy for arguing for some *subjective constraint* on obligation—if what we're able to do is in some sense constrained by our subjective perspective, then, given OIC, our subjective perspective must similarly constrain our obligations.¹

Something like this very general thread can be found across a variety of metaethical accounts—Bernard Williams, Christine Korsgaard, Richard Joyce, and, more recently, Errol Lord

¹ Two points of clarification: 1) by 'subjective perspective', I just mean how the world seems to be—normatively and non-normatively—from a particular agent's point of view; 2) I use 'ought' and 'obligated' interchangeably and construe them practically; I have no special moral sense in mind.

can all be read as offering variations on this theme.² But while such a strategy appears attractive, it's fraught with pitfalls. Ability comes in many senses, and the truth of ability ascriptions is (at least very often) context sensitive. Given this, I worry that attempts to draw metaethical conclusions from substantive claims about ability trade on the plausibility of ability ascriptions without taking sufficient care to make explicit the operative sense of ability operating throughout the argument.

Naturally, this lack of precision causes problems. But I don't take this to show that appeals to ability are destined to be metaethically fruitless. On the contrary, my hope is that getting clear on the candidate senses of ability that obligation might entail will suggest an underexplored line of philosophical inquiry.

In what follows, I begin by highlighting the trouble that the context sensitivity of ability ascriptions makes for metaethical debates by examining Lord's recent attempt to argue for an *epistemic* constraint on obligation from premises that appeal to agents' abilities.³ In so doing, I make explicit how different theories of obligation—specifically, Objectivism, Constructivism, and Lord's Perspectivalism—take obligation to be constrained by different *senses* of ability. Given these competing senses, if an appeal to ability (like OIC) is to serve as a starting point for a metaethical debate, it must be supplemented with a defense of that sense of ability as the normatively relevant one. I conclude with a preliminary sketch of just such an argument.

LORD'S PERSPECTIVALISM

Lord argues that what an agent is obligated to do is a function of the reasons to which she has epistemic access.

Perspectivalism: what an agent is obligated to do is a function of only the reasons to which she bears the right epistemic relation.⁴

The motivating thought behind Perspectivalism is that "the facts that obligate must be potentially action-guiding in a certain sense—[they] must at least potentially be the reasons *for which* we act." And, Lord insists, "we can have the ability to act for the right reasons only if we possess those reasons," where 'possessing' those reasons amounts to bearing the right epistemic relation to them.⁵

Lord contrasts his Perspectivalism with *Objectivism* about obligation.

Objectivism: what an agent is obligated to do is a function of all of her reasons.⁶

² Williams (1981); Korsgaard (1986); Joyce (2001); Lord (2015). Notably, Jonny Anomaly (2008) explicitly argues that we read Williams in this way, and Chris Heathwood (2011) suggests a similar reading of Korsgaard. I'm much obliged to Alex Hyun for pointing me towards many of these sources, and for his thoughtful discussion of them.

³ Lord (2015).

⁴ In a footnote, Lord mentions that his considered view is that the right epistemic relation is *being in a position to know*. But none of his arguments commit him to anything so specific. Lord (2015): 29, fn 5.

⁵ Lord (2015): 27.

⁶ *Ibid*: 28.

Objectivism insists that an agent's obligations are a function of the facts that speak for and against her actions, whether she knows them or not, whether they motivate her or don't.⁷ Given this, it's easy to see why Perspectivalism and Objectivism conflict at the most basic level—the Perspectivalist insists while the Objectivist denies that there is a subjective constraint on obligation.

Importantly, though, Objectivism is not Perspectivalism's only foil (though it's the only one Lord considers). Take Sharron Street's metaethical Constructivism:

Constructivism: what an agent is obligated to do is a function of the agent's practical point of view together with the non-normative facts.⁸

Constructivists of Street's variety agree with the Perspectivalist that there is some subjective constraint on obligation, but disagree about the nature of the constraint. The Perspectivalist construes the constraint epistemically; the Constructivist construes it motivationally. So Lord's project shouldn't be understood as simply a critique of Objectivism, defending some generic subjective constraint on obligation.⁹ Rather, Perspectivalism defends a specific *epistemic* constraint. This is important because, as we'll see in the next section, Lord's assessment of ability is dubious not just in the eyes of the Objectivist, but also in the eyes of those (like the Constructivist) who endorse conflicting subjective constraints.

In defense of Perspectivalism, Lord offers the following argument:

1. **Right Reasons Ability Condition (RRAC):** "If *A* ought to φ , then *A* has the ability to φ for the right reasons."
2. **Possession:** "If *A* has the ability to φ for the right reasons, then *A* possesses the right reasons."
3. So, "If *A* ought to φ , then *A* possess the right reasons."¹⁰

Lord's first premise is a particular, somewhat more contentious instance of the familiar thought that ought implies can. While RRAC is not as obviously uncontroversial as OIC, I suspect that most metaethicists would be inclined to endorse something like it.¹¹ But that doesn't stop Lord from providing a robust defense. First, he argues that unless RRAC is true, "there will be cases where one ought to φ even though one is unable to φ non-accidentally in the right way."¹² The thought is that if RRAC is false, then an agent might be obligated to do something that she could

⁷ Though Lord does not specifically identify anyone as an Objectivist, Derek Parfit (2011) fits the bill.

⁸ Street (2010).

⁹ To be fair to Lord, while he only mentions the Objectivist as an opponent, he never claims that there are no others.

¹⁰ Lord (2015): 35. "Right Reasons Ability Condition" is Lord's Label. "Possession" is my label. Again 'possession' is just short hand for 'bears the right epistemic relation to'.

¹¹ Joseph Raz (2011) is at least one notable exception. He argues that agents can have "non-standard" reasons for action that cannot be (directly) acted upon. Idiosyncratic views aside, though, even the Objectivist—Lord's chief opponent—can accept RRAC given an expansive enough understanding of ability. Lord agrees: "[objectivists] feel no need to deny the Right Reasons Ability Condition [...] as long as they hold a liberal view of what it takes to have the ability to act for the right reasons." (2015): 38. As we'll see, though, this is the very thing that causes problems for his strategy and others like it.

¹² Lord (2015): 36.

only do unintentionally, by mistake, or for some reasons other than those which genuinely call for the action. But this isn't how obligation works—our obligations are obligations to *respond* to the reasons that favor the obligated actions. So we must be able to act for the reasons that determine our obligation. Second, Lord argues that unless RRAC is true, "there will be cases where one ought to ϕ even though one is unable to ϕ in a way that would be creditworthy."¹³ Why? Because one deserves credit for doing what she ought only if she does what she ought for the reason(s) she ought. And if an agent can be obligated to do something despite being unable to do it for the right reason (i.e., if RRAC is false), then she might be obligated to do a thing that she could not do in a creditworthy way. But doing what one ought to do is *always* creditworthy—fulfilling an obligation is an achievement, after all!—so RRAC must be true.

The second premise of Lord's argument insists that the ability to act for the right reasons requires that an agent be in the right epistemic contact with those reasons.¹⁴ In defense of this, Lord asks us to consider a pair of cases.

Delusional Andy (D-Andy): "Andy knows that his wife has always been an extremely loyal person. He also knows that he has no reason to think that she is cheating on him. Despite this knowledge, he does believe that she is cheating on him. He thus files for divorce. In fact, his wife is cheating on him."

Surprised Andy (S-Andy): "Andy knows that his wife has always been an extremely loyal person. However, much to his surprise, he learns that she is cheating on him—her best friend tells him, he finds some love letters, and he catches his wife with her lover. He thus files for divorce."¹⁵

Intuitively, S-Andy acts for the right reasons, but D-Andy does not. After all, S-Andy (in some sense) *grasps* the right reasons in a way that D-Andy does not. Moreover, if we insist (against intuition) that D-Andy acts for the right reasons, then, Lord argues, we're committed to thinking (i) that D-Andy's act is creditworthy (because he does what he ought for the reason that he ought); (ii) that he is sensitive to the reasons for and against his action (because acting for the right reason entails being responsive to the relevant considerations); and (iii) that his action is justified (because doing what one ought for the reason that one ought explains why that action is justified).¹⁶ But none of this seems right—D-Andy doesn't deserve credit, he isn't sensitive to his reasons, and his action isn't justified. All this confirms the intuitive judgment—only S-Andy acts for the right reason.

But what *explains* this? Why does S-Andy act for the right reason while D-Andy doesn't, given that they perform identical actions and are motivated by identical considerations? The only difference between them is that S-Andy *knows* his wife is cheating on him while D-Andy doesn't. And this seems like a plausible explanans—if reasons are *facts* (as most philosophers suppose), then acting for the right reason plausibly requires some connection to the fact that is the right

¹³ *Ibid.*: 37.

¹⁴ Again, Lord's considered view is that the right sort of epistemic contact is *being in a position to know*.

¹⁵ *Ibid.*: 38.

¹⁶ *Ibid.*: 38-40. Lord offers a more robust defense of each of these implications than there is space to cover here. Each is, I think, intuitive, if a bit opaque.

reason. But D-Andy doesn't have *access* to the fact that is the right reason for him to get a divorce. He is *motivated* by a belief whose content *happens to be* the right reason for him to get a divorce, but only because he got "lucky" in the sense that his delusion turned out to be true.¹⁷ Moreover, D-Andy doesn't even seem *capable* of doing of acting for the right reasons—given his epistemic lot, there's just no way for him to apprehend the right reason in a way that would allow him to genuinely act for it. All this speaks in favor of Possession—an agent has the ability to ϕ for the right reasons only if she possesses those reasons.

So goes Lord's defense of RRAC and Possession. From them, Perspectivalism seems to follow. Again, Lord isn't the only one who pursues this kind of strategy, but his argument is an especially clear instance of it. In the next section, I explain my worry with it—why I think that it equivocates over multiple senses of ability.

WHAT MAKES ANDY ABLE?

Lord's argument for Perspectivalism is persuasive. But conspicuously absent from it is any discussion of what *ability* amounts to.¹⁸ And this is worrisome. While both premises look plausible, and are well defended, there's no consensus on the right way to analyze ability.¹⁹ Even more worryingly, the truth of ability ascriptions appears, in at least one important respect, to be (almost always) context sensitive. To illustrate, take the following ability claims.

Ab₁ - Katie is unable to leap the canyon.

Ab₂ - Clara is unable to look her mother in the eye.

Ab₃ - Luke is unable to attend the party.

Suppose all three of these claims are true—the canyon is too wide for Katie, Clara is too embarrassed to face her mom, and Luke already committed to babysitting his niece. The crucial thing to note is that what *makes* each claim true is different in each case. In Ab₁, Katie's inability must be explained by something like the fact that there is no possible world in which she (i) is in (more or less) identical circumstances, (ii) has (more or less) identical physical traits, and (iii) actually leaps the canyon.²⁰ (Similarly, her *ability*, if she had it, would be explained by the *existence* of such a world). In Ab₂, Clara's inability can't be similarly explained—clearly there exists a world in which she is in (more or less) identical circumstances, with (more or less) identical physical traits, and looks her mother in the eye. Instead, Clara's inability must be explained by something like the lack of a world in which she looks her mother in the eye *given* her external circumstances *and her current psychological makeup*. In Ab₃, Luke's inability demands a different explanation still—his must be explained by the lack of a world in which he attends the party *given* his external circumstances *and that he fulfills his conflicting obligations*. So while each of these ability claims

¹⁷ Epistemically, lucky, of course. The fact that his wife is cheating on him is decidedly unlucky.

¹⁸ Lord mentions (in a footnote) that he takes abilities to be a kind of disposition. He also mentions (in a separate footnote) that he doesn't take dispositions to be analyzable counterfactually. I take the worries expressed here to apply despite these clarifications. Lord (2015): 34, fn 11; 40, fn 20.

¹⁹ For a helpful overview of the competing positions, see Maier (2014).

²⁰ By physical traits/circumstances, I mean *merely* physical—i.e., physical, but not mental, epistemic, or normative.

might be true, each has different truth-makers. To secure the truth of each ascription, each sentence implicitly makes reference to a different set of relevant possibilities.

Angelika Kratzer convincingly argues that ability claims are always relational in this way—they relate some event (leaping the canyon, looking in the eye, attending the party) to some restricted set of worlds.²¹ The clause that restricts the set of worlds is just as essential to the meaning of an ability ascription as is the specification of the event. But the restriction is often implicit (as in Ab_{1-3}), and when it's implicit, it must be derived from context. Usually this is no trouble. For example, Ab_{1-3} are naturally understood as implicitly expressing the following *given* clauses:

Ab_1^* - Katie is unable to leap the canyon [*given her physical traits*].

Ab_2^* - Clara is unable to look her mother in the eye [*given her psychological state*].

Ab_3^* - Luke is unable to attend the party [*given his prior commitments*].

Call the *given*-clause the *restrictor*. An ability ascription, then, is true just in case the agent performs the relevant action in at least one of the restricted set of worlds.²² The essential point is that we cannot determine the truth of an ability ascription without identifying the relevant restrictor. And while it's usually harmless to assume a shared understanding of the restrictor in everyday language, it might not be so innocuous in the context of a metaethical dispute.²³

Once we appreciate that the truth of an ability ascription depends on an often implicit characterization of the relevant possibility space, we should wonder what restriction is operative in the ability claims of RRAC (and OIC) and Possession.

Focus on Possession first. What makes S-Andy able and D-Andy unable to act for the right reason? Presumably, it's the feature that Lord identifies—S-Andy is able, while D-Andy is unable, *given* what each has epistemic access to. In other words, S-Andy is able to act for the right reason because there exists some world in which S-Andy (i) is in (more or less) identical circumstances, (ii) has epistemic access to (more or less) the same things that he does in the actual world, and (iii) acts for the right reason. D-Andy, on the other hand, lacks the ability to act for the right reason because there's no world in which he acts for the right reason *given* his external circumstances and

²¹ Kratzer (1977). Kratzer's discussion is in terms of "can," rather than "able", but all the same arguments apply. I should also mention that it's controversial whether abilities are appropriately analyzed in terms of possible worlds. Perhaps they're not. Regardless, I strongly suspect that something like this context sensitivity will arise on any plausible analysis of ability. Arguing this point, though, is beyond my abilities here.

²² Or perhaps in one of the "closest" of these worlds. But these details needn't concern us here. See Peacocke (1999): ch. 7.3 for a discussion of the relevance of closeness in ability ascriptions.

²³ It's worth mentioning here that, for ease of explication, I've been fudging one of Kratzer's more important insights. It's misleading to speak, as I have been, in terms of different *senses* of ability. There are just too many ways in which we can plausibly restrict the possibility space of an ability ascription to suppose that there's a distinct *sense* of ability for each restriction. What Kratzer shows is that ability is *univocal* in the sense that it always does the same thing in an expression—it relates some event to some restricted set of worlds. The variable that *changes* (again, often implicitly, with context) is the *restriction*. So '*ability*'s meaning is static. It's the relata of ability ascriptions are variable. With this caveat in mind—again, for ease of explication—I'll continue to speak in terms of different "senses" of ability.

epistemic profile. So the implicit restrictor in Possession limits the possibility space to externally and *epistemically* similar worlds. For convenience, let's call so-restricted ability *P-ability*.

P-ability: An agent is P-able to φ just in case there is some possible world in which the agent's physical and epistemic circumstances are (more or less) the same, and that agent intentionally φ -s.²⁴

Crucially, understanding the Andys' abilities in something like this way is *necessary* to make Possession plausible. Clearly it's not the case that D-Andy lacks the ability to act for the right reason given *just* his external circumstances and physical traits. Also importantly, this isn't yet a criticism of Possession. The sense of ability required to make Possession plausible—P-ability—is a perfectly respectable one. The point thus far is just that not any old restrictor will do.

The *problem* arises when we export P-ability to RRAC. Above, I mentioned that RRAC is (relatively) uncontroversial on at least *some* sense of ability. Very plausibly, when an agent is obligated to φ , there must be *some* possible world in which she successfully φ -s for the right reason. But P-ability isn't nearly so weak—the P-ability to act for the right reason requires that an agent act for the right reason in some world in which her *epistemic* circumstances are (more or less) the same. But an epistemically restricted sense of ability is one that Lord's opponents will simply reject when it's placed in RRAC. The Objectivist can (and should, I think) admit that P-ability is a perfectly respectable sense of ability, and that Possession is true *given* that sense of ability, all while *denying* that *obligation* entails a similarly restricted ability to act for the right reason. Instead, the Objectivist can consistently maintain that while Possession is true *given P-ability*, *obligation* entails the *O-ability* to act for the right reason.

O-ability: an agent is O-able to φ just in case there is some possible world in which the agent's physical circumstances are (more or less) the same, and that agent intentionally φ -s.

To secure the ability to act for the right reason, O-ability places no psychological or epistemic restriction on the set of possible worlds in which the agent successfully acts for the right reason. Similarly, the Constructivist can grant Possession *given P-ability*, while consistently insisting that *obligation* entails something else—the *C-ability* to act for the right reason.

C-ability: an agent is C-able to φ just in case there is some possible world in which the agent's physical and motivational circumstances are (more or less) the same, and that agent intentionally φ -s.

Like the Objectivist, the Constructivist can simply deny that the sense of ability in RRAC has anything do to with the obligated agent's epistemic lot.

It's worth emphasizing here that the Objectivist and Constructivist might make these moves *while still accepting* each of Lord's premises for all of the reasons he endorses them; they just insist

²⁴ Two things:

1. The standard view seems to be that obligation entails the ability to do things *intentionally*. I agree with this, so I included it. But defending this claim is beyond the scope of this paper. If you object to the inclusion of "intentionally", feel free to ignore it.
2. Let an agent's physical circumstances include her physical traits.

that a different sense of ability is operative in each. The point here isn't that Lord's argument is incoherent—Perspectivalism has a perfectly respectable story to tell about ability, and Possession and RRAC are plausible enough with P-ability plugged in. Rather, the worry is that Lord's argument is *dialectically* ineffective. The sense of ability required to make Possession plausible is a sense of ability that, when plugged into RRAC, makes RRAC easy for Lord's opponents to deny.

Given this, despite first appearances, Lord's argument doesn't seem very well-suited to advance the metaethical debate. And I think this point generalizes. While most metaethicists are happy to endorse some version of RRAC (or OIC), it doesn't represent the point of consensus that it seems to. This is because different metaethical approaches take different and *inconsistent* senses of ability to constrain obligation. Consequently, unsupplemented appeals to ability are more likely to obscure rather than illuminate genuine sources of metaethical disagreement. So long as each party to a metaethical dispute has an internally consistent story to tell about the normatively relevant sense of ability, substantive ability claims will have trouble gaining traction absent some explanation of why *that sense* of ability is the normatively relevant one.

SO WHAT CAN ABILITY DO?

Does all this mean that thinking about ability is metaethically useless? I don't think so. But I take the lesson of the forgoing discussion to be that serious problems arise when appeals to ability lay at the *foundation* of metaethical debates. The context sensitivity of ability ascriptions makes it too easy to obscure genuine points of contention with agreeable-*sounding* ability claims. The result is that disputing parties end up talking past one another. The solution, I think, is not to abandon ability talk all together, but to offer a *defense* of a particular sense of ability *as the normatively relevant one*. In this section, I make a preliminary attempt to do so—arguing that only an *epistemically and motivationally* restricted sense of ability can capture the full extent to which an agent's options are limited by her circumstances.

The general strategy is this: take a certain characterization of ability and ask whether all the actions it categorizes as able-to-be-performed could serve a good answer to the central deliberative question—'*what should I do?*' Obligations pick out the unique answers (when they exist) to the central deliberative question.²⁵ So, if a sense of ability characterizes an agent as-able-to-perform some action that *is ineligible* to answer the central deliberative question, then we've good reason to think that that sense of ability isn't restrictive enough.

To see this strategy in action, consider the following two famous cases:

Miners. A group of 10 miners are trapped in a mine. They are either trapped in shaft A or in shaft B. It is not easily knowable which shaft they are in. Flood waters are approaching the shafts. Billy has the choice to sandbag shaft A, sandbag shaft B, or not sandbag either. She knows that if she sandbags A and the miners are in A, all the miners will survive. She knows the same is true of B. She also knows that if she sandbags either shaft and the miners are in

²⁵ The "central deliberative question" is how Lord (2015) isolates the sense of obligation of interest to his project. I'm indebted. I find it to be an extremely helpful way to pick out the "all-things-considered ought", and it serves my purposes here quite well.

the other shaft, they will all die. Finally, she knows that if she does nothing, then 9 of the 10 will survive. What should she do?²⁶

Squash. Stan has lost a very hard game of squash to an infuriating opponent. Stan knows that, in this situation, the virtuous person would calmly shake his opponent's hand. But he also knows that if *he* were to attempt that, he would fly into a rage and hit his opponent with his racquet. What should he do?²⁷

Take Miners first. We might recommend that Billy sandbag whichever shaft the miners are in. But to this, she might reasonably reply: "I would love to, if I only knew which shaft that was. But I don't. So I can't. So *what should I do?*" The fact that Billy can still ask the central deliberative question suggests that the action we offered failed to answer it. So blocking the miners' shaft can't be the action that Billy is obligated to perform. Moreover, Billy's explanation for *why* our recommendation was bad seems right—she isn't obligated to block the miners' shaft (i.e., "block the miners' shaft" is not a viable answer to the central deliberative question) because Billy isn't *able* (in the relevant sense) to do so. But clearly Billy is O-able and C-able to sandbag the miners' shaft—there exists some possible world with (more or less) identical physical and motivational circumstances in which Billy sandbags the shaft which contains the miners. So O- and C-ability don't capture the extent to which our obligations are constrained by our abilities.

Does P-ability do better? It seems to fare well in Miners.²⁸ It correctly categorizes Billy as unable to sandbag the shaft the miners are in—she isn't P-able to sandbag the miners' shaft because there's no possible world in which she is in (more or less) the same epistemic circumstances and intentionally does so.²⁹ But P-ability doesn't seem to get Squash right. We might recommend to Stan that he go congratulate his opponent. But to this he might reasonably reply: "Surely that would be best, and I would love to do so. But I *can't*. I just know that I'll beat him with my racquet. So *what should I do?*" Again, we've failed to answer the central deliberative question. And, plausibly, the *reason* we've failed to answer the central deliberative question is that we've cited an action which Stan is not *able* (in the relevant sense) to perform. But clearly Stan is P-able (and O-able, incidentally) to shake his opponents hand (there exists some possible world with (more or less) identical physical and epistemic circumstances in which Stan shakes his opponents hand). So P-ability doesn't capture the extent to which our obligations are constrained by our abilities.

It would be too quick to draw any firm conclusions from this. But it at least suggests, I think, that our obligations are constrained by some sense of ability that is *both* epistemically *and* motivationally restricted. Only a dually restricted sense of ability could get *both* Miners and Squash right. Precisely how we should characterize these constraints is beyond my purposes

²⁶ This is Lord's (2015) statement of the case. It was conceived by Donald Regan (1980) and made famous by Derek Parfit (1988).

²⁷ This is Kieran Setiya's (2007) statement of the case (with some slight, non-substantive modifications). It was conceived by Gary Watson (1975).

²⁸ This isn't surprising. Miners is one of the cases Lord uses to motivate Perspectivalism.

²⁹ To intentionally sandbag the shaft the miners are in, she would need a justified belief about which shaft they were in.

here. My more modest goal is to gesture towards a promising way to argue for a particular sense of ability as the normatively relevant one. Generally, the sense of ability that constrains obligation should capture the extent to which an agent (*qua* agent) has limited control over the deliberative problems she faces. To the extent that we can get an independent grip on the bounds of agency, we can appeal to such considerations to hone in on the normatively relevant sense of ability. And once we do this, given that ought implies can, ability might be able to do some metaethical work after all.

BIBLIOGRAPHY

- Anomaly, Jonny (2008). Internal Reasons and the Ought-Implies-Can Principle. *Philosophical Forum* 39 (4): 469-483.
- Heathwood, Chris (2011). Desire-Based Theories of Reasons, Pleasure, and Welfare. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics, vol. 6*. Oxford University Press: 79-106.
- Joyce, Richard (2001). *The Myth of Morality*. Cambridge University Press.
- Korsgaard, Christine (1986). Skepticism about Practical Reason. *Journal of Philosophy* 83 (1):5-25.
- Kratzer, Angelika (1977). What 'Must' and 'Can' Must and Can Mean. *Linguistics and Philosophy* 1 (3): 337-355.
- Lord, Errol (2015). Acting for the Right Reasons, Abilities, and Obligation. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics, vol. 10*. Oxford University Press: 26-52.
- Maier, John (2014). Abilities. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. URL = <<http://plato.stanford.edu/archives/fall2014/entries/abilities/>>.
- Parfit, Derek (2011). *On What Matters*. Oxford University Press.
- Parfit, Derek (1988). What We Together Do, unpublished
- Peacocke, Christopher (1999). *Being Known*. Oxford University Press.
- Raz, Joseph (2011). *From Normativity to Responsibility*. Oxford University Press: Ch. 3.
- Regan, Donald (1980). *Utilitarianism and Cooperation*. Oxford University Press: 265.
- Setiya, Kieran (2007). *Reasons without Rationalism*. Princeton University Press.
- Street, Sharon (2010). What is Constructivism in Ethics and Metaethics? *Philosophy Compass* 5 (5): 363-384.
- Streumer, Bart (forthcoming). Reasons and Ability. In Daniel Star (ed.) *The Oxford Handbook of Reasons and Normativity*. Oxford University Press.
- Vranas, Peter B. M. (2007). I Ought, Therefore I Can. *Philosophical Studies* 136 (2): 167-216.
- Watson, Gary (1975). Free Agency. *Journal of Philosophy* 72 (April): 205-220.
- Wedgwood, Ralph (2013). Rational 'Ought' Implies 'Can'. *Philosophical Issues* 23 (1): 70-92.
- Williams, Bernard (1981). Internal and External Reasons. In *Moral Luck*. Cambridge University Press: 101-113.