

WHAT ABILITY CAN DO¹

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 maintain that while such a strategy appears promising, it is fraught with pitfalls. Specifically, I argue that because the truth of an ability ascription depends on an (almost always 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, appeals to ability in metaethics aren't doomed to be fruitless. On the contrary, if we can independently establish a particular sense of ability as the normatively relevant one, then we have good grounds for ruling out metaethical accounts that are inconsistent with it. In the final section, I make just such an argument. What seems right about the thought that ability constrains obligation is that an agent cannot be obligated to do that which her *circumstances* prevent her from doing. I argue that only a sense of ability that is both epistemically and motivationally restricted adequately respects the limits of agential control.

I. INTRODUCTION

Here's a natural thought: what we're obligated (or have reason) to do is constrained by what we're able to do.² This idea is both intuitive, and widely endorsed by philosophers of otherwise disparate metaethical bents. As such, it appears to be a solid starting point for a metaethical debate. In particular, it suggests a promising strategy for arguing for the existence of some subjective constraint on obligation—if an agent's subjective perspective limits what she is able to do, and what she is able to do limits what she is obligated to do, then her subjective

¹ Penultimate draft. Final version published in *Philosophical Studies* (<https://link.springer.com/article/10.1007/s11098-017-0888-3>).

² I suspect that our abilities constrain *both* our obligations and our practical reasons. But in what follows, I'll primarily discuss the prospect of our *obligations* being so-constrained. I do this for two reasons: (i) the idea that ability constrains obligation is less contentious than the analogous claim about reasons, and (ii) the metaethical accounts I discuss are primarily concerned with obligations.

perspective constrains her obligations.³ But while such a strategy appears attractive, I'll argue that it is fraught with pitfalls. Ability comes in many senses, and the truth of an ability ascription is (at least very often) context sensitive.⁴ Given this, I worry that attempts to draw metaethical conclusions from substantive claims about ability often trade on the plausibility of ability ascriptions without taking sufficient care to make explicit the 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 open up an underexplored line of philosophical inquiry. Specifically, if the sense of ability that's entailed by obligation is one that's supposed to capture the extent to which an agent exercises genuine control over her circumstances (as I think is plausible), then we can draw on the resources of the philosophy of action and agency to determine precisely which sense of ability respects the limits of this control. Ultimately, I'll suggest that only a sense of ability that is epistemically and motivationally restricted is capable of doing the job. This, in turn, suggests that our obligations are similarly subjectively constrained.

In what follows, I begin by getting clear about the sense of obligation I have in mind, and offer some brief considerations in support of the thought that our abilities constrain our obligations. I then sketch what I take to be an attractive strategy for arguing for a subjective constraint on obligation from premises that appeal to agents' abilities. Next, I highlight the trouble that the context sensitivity of ability makes for metaethical debates by examining a particular instance of the attractive strategy—Errol Lord's (2015) recent argument for an epistemic

³ Two points of clarification: (i) by 'subjective perspective', I just mean how the world seems to be—normatively and non-normatively—from a particular agent's point of view; (ii) I use 'ought' and 'obligated' interchangeably and construe them practically (I'll say a bit more about this in the next section).

⁴ As we'll see, it's a bit misleading to speak of different *senses* of ability. Rather, the point is that an ability ascription relates a particular action to a restricted set of possible worlds, and it's the character of the *restriction* rather than the sense of ability that's context sensitive. That said, it's far easier to talk in terms of senses of ability. So I'll continue to speak this way throughout.

constraint on obligation. In so doing, I try to make explicit the precise sense in which three different metaethical theories—Objectivism, Constructivism, and Perspectivalism—take ability to constrain obligation. I then show how each view carries different commitments about the boundaries of agency—i.e. the extent to which we exercise genuine agential control over our circumstances. I conclude by suggesting that in order to respect the limits of agency, we must recognize both an epistemic and motivational constraint on obligation.

II. OUGHTS AND ABILITIES

Different philosophers (and different folk, for that matter) mean different things when speaking of oughts and obligations. So it's important to be clear at the outset about the sense that's relevant for our purposes here. First, I make no distinction between what an agent ought to do and what she is obligated to do; I use 'ought' and 'obligated' interchangeably. Second, I'm focused on *all-things-considered, practical* obligation, not epistemic, moral, or any other sort.⁵ Finally, I'm after a sense of obligation that is essentially *normative*—a sense that speaks to what an agent *should do*.⁶ There are, admittedly, a number of dimensions along which we might evaluate an action. An action might constitute (or partially constitute) an auspicious event, such that the world is better for its being performed. An action might manifest a virtuous disposition, such that the agent who performs it (or tends to perform it) is of higher character than someone who does not. Plausibly, an action might be evaluated in a host of other ways. But none of these evaluations will be genuinely normative (in the sense that I'm after) unless they help explain why the action is *to be done*, or why the action is good *as an action*. When an agent does what she ought to do, she succeeds *in acting*; when she does not, it's her *action* (perhaps among other things) that is defective. All this, on its face

⁵ For what it's worth, to the extent that other kinds of oughts (epistemic, moral, prudential, etc.) *are* constrained by abilities, I suspect that a similar project could be mounted in terms of each. But I'm much less sure that other kinds of oughts are so-constrained. And even if they are, the way in which they are might be importantly different from the all-things-considered, practical ought. So for this project, I set them aside.

⁶ As distinguished from *evaluative*. See Wedgwood (2009) for a particularly insightful discussion of the normative and evaluative domains, and their relation to one another.

anyway, is independent of whether the resulting state of affairs is good or bad, or the acting agent is virtuous or vicious, etc.

The most perspicuous way to isolate this sense of obligation, I think, is as Lord (2015, p. 28) does: an agent's obligation picks out the unique answer (when there is one) to the "central deliberative question"—*What should I do?* If the unique answer to the central deliberative question is ϕ , then the inquiring agent ought to ϕ .^{7,8}

Given this characterization of obligation, are our obligations constrained by our abilities? Surely they must be in some sense. Take two similar cases:

Anne: Sitting in a lakeside park, Anne spots a boy who has fallen out of his canoe and is struggling to stay afloat. He appears to be drowning. Anne is a strong swimmer and would have no problem reaching the boy in time to save him. What should she do?

Beth: Sitting in a lakeside park, Beth spots a boy who has fallen out of his canoe and is struggling to stay afloat. He appears to be drowning. Beth never learned how to swim. She knows that if she were to dive into the lake, not only would she be of no help, but two people would then need saving. What should she do?

Clearly, Anne should dive in and save the boy. Just as clearly, I think, Beth shouldn't. There may, of course, be other things Beth should do. But she shouldn't dive in the lake. Since the only apparent difference between Anne and Beth is that Anne has an ability (to dive in and rescue the boy) that Beth lacks, it's reasonable to suppose that it's this difference in ability that explains the difference in their respective obligations.⁹

⁷ Naturally enough, when there are multiple correct answers to the central deliberative question, those answers pick out the actions that would be *permissible*, but not obligatory, for the inquiring agent to perform.

⁸ What's the connection between this sense of obligation and practical reasons? I'm inclined to think that what an agent is obligated to do just is what she has most practical reason to do. But there are different views about how reasons relate to obligations. Lord (2015), for example, takes an agent's obligations to be determined by a *subset* of her reasons. Streiffer (2003) takes a similar approach. Gert (1998: ch. 3) argues for a more complicated relationship between reasons and obligations. To avoid this controversy, I'll largely avoid speaking in terms of reasons (though at times, my discussion of Lord will force my hand).

⁹ Of course none of this is to deny that it would be *better* for Beth to dive in and (successfully) save the boy. Clearly that would be ideal. But given her inability to do so, that option isn't in the cards. And so, at least in this case (and, I'm inclined to think, in virtually all cases), the answer to the "what would be best?" question comes apart from the answer to the central deliberative question ("what should she *do*?").

At the risk of hastily generalizing, I'm inclined to hastily generalize: cases like these show that *in general* our abilities (somehow) constrain our obligations. Put more pithily (and more familiarly), ought implies can.

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

As mentioned at the outset, this isn't a particularly controversial thought as far as philosophical claims go—that ought implies can has been explicitly defended by a number of philosophers, and taken for granted by many others.¹⁰ Of course, even among those who agree that ability somehow constrains obligation, opinions vary about the precise character of the ability-constraint. For my purposes here, I'll take OIC to be an umbrella principle, and consider any principle that posits some ability-constraint on obligation (or practical reason) to be a *version* of OIC.¹¹ Naturally there are some philosophers who deny that ability constrains obligation in any sense whatsoever.¹² But (again, as far as philosophical claims go) OIC, in its most general form, enjoys broad support.

Given that OIC is relatively uncontroversial, it appears to be an attractive starting point for a metaethical debate—if abilities (somehow or another) 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 sort of *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. Here's the argument, in very general terms:

¹⁰ A small sampling of explicit defenses: Vranas (2007) offers what I take to be a convincing defense of OIC; Streumer (forthcoming) defends the related thesis that reasons imply can; Wedgwood (2013) defends a rational-ought-implies-can principle.

¹¹ This way of thinking about OIC is, admittedly, a bit idiosyncratic. Advocates of more specific ability-constraints on obligation might object to their principles being considered "*versions* of OIC". I speak this way simply out of convenience—to readily pick out the class of principles that take ability to (somehow) constrain obligation. In so doing, I don't mean to imply that all versions of OIC are equally plausible or similarly supported.

¹² In his defense of OIC, Vranas (2007) surveys wide range of counterexamples and objections. Graham (2011) offers a more recent, forceful critique of the principle. But engaging with criticisms of OIC is beyond my purposes here. My intention is just to briefly motivate OIC, and then take it for granted in what follows.

An Attractive Strategy

1. **Ought Implies Can (OIC):** If an agent ought to φ , then she is able to φ .
2. **Subjectively Constrained Ability (SCA):** An agent's ability to φ is constrained by her subjective perspective.
3. Therefore, an agent's subjective perspective constrains her obligations.

Something like this general thread can be traced through a variety of metaethical accounts—Bernard Williams (1981), Christine Korsgaard (1986), Richard Joyce (2001), Stephen Darwall (1992), Michael Huemer (2005), and (most relevantly for the current project) Lord (2015) can all be read as offering variations on this theme. Unsurprisingly, these philosophers employ the attractive strategy for diverse purposes, and the details vary—about precisely how ability constrains obligation and the extent to which an agent's abilities are limited by her subjective perspective—but each can be read as either assuming or defending some version of OIC, maintaining that an agent's subjective point of view limits what that agent is able to do, and concluding that an agent's obligations (or reasons) are constrained by her subjective perspective.¹³

Though the strategy is attractive, I worry that to the extent that it depends on substantive claims about agents' abilities, it masks deep disagreement over what *sense* of ability is the normatively relevant one. If this is right, then what seemed to be a shared starting point may actually represent *conflicting* commitments about the relationship between obligation and ability—the sense of ability an Objectivist has in mind when reading OIC might be very different from the sense of ability operating in a Subjectivist's understanding of what is ostensibly the same (or at least a similar) principle.

The best way to appreciate this problem, I think, is to look carefully at a version of the attractive strategy in action. In the next section, I examine Lord's argument for a subjective (epistemic) constraint on obligation from premises appealing to agents' abilities. I'll attempt to

¹³ I'll have a bit more to say about these philosophers' use of the attractive strategy in section III. I'm very thankful to Alex Hyun for pointing me towards many of these sources, and for his thoughtful discussion of them.

show that because the truth of ability ascriptions is context sensitive, his argument subtly equivocates between two senses of ability. I single out Lord for two reasons: first, his argument is an especially (and refreshingly) clear instance of the attractive strategy; second, but for the problem I highlight, I find his argument to be quite persuasive. Given these, Lord's argument serves as a particularly good example of how an appeal to ability can obscure deep metaethical disagreement. It's important to keep in mind throughout, though, that I don't take the worries I raise to apply uniquely to Lord. Rather, I take the problem to apply generally to any attempt to draw metaethical conclusions from a version of OIC without substantiating the operative sense of ability.

III. THE ATTRACTIVE STRATEGY IN ACTION – LORD'S PERSPECTIVALISM

Before examining Lord's argument, it will be helpful to have his view and its competitors on the table. Lord argues for an *epistemic* constraint on obligation—maintaining that what an agent is obligated to do is a function of the reasons to which she has the right kind of epistemic access.¹⁴ He calls his view *Perspectivalism*.

Perspectivalism: what an agent is obligated to do is a function solely of the reasons to which she bears the right epistemic relation.¹⁵

According to Lord (2015, p. 27), 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, as we'll see shortly, he maintains that "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 those reasons.

¹⁴ By 'reasons', Lord (2015, p. 28) means "facts that recommend certain reactions." He doesn't take there to be any subjective constraint on what counts as a reason. So, on his view, what an agent is obligated to do is determined by a *subset* of her reasons. Again, this is not how I'm inclined to understand the relationship between reasons and obligations. But that's not my beef here. So I do my best to ignore it in what follows.

¹⁵ In a footnote, Lord (2015, p. 29, fn 5) 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 contrasts his Perspectivalism with *Objectivism* about obligation.

Objectivism: what an agent is obligated to do is a function of all of her reasons.¹⁶

As Lord characterizes the Objectivist, she denies that there is *any* subjective constraint on what agents are obligated to do. According to the Objectivist, 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. 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).¹⁷ Take Sharon Street's (2010) metaethical Constructivism:

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

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. 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 who (like the Constructivist) endorse conflicting subjective constraints.¹⁸

¹⁶ Lord cites Moore (1912), Thomson (1986), and Graham (2010) as advocates of Objectivism, so construed.

¹⁷ To be fair to Lord, while he only mentions the Objectivist as an opponent, he never claims that there are no others.

¹⁸ Constructivism isn't the only non-Objectivist competitor. There are other accounts of obligation that conflict with Lord's view. I focus on Constructivism only because I take the disagreement between the Constructivist and Perspectivalist to be particularly easy to pinpoint. In addition, various accounts of practical reasons might intersect with or mirror these accounts of obligation in interesting ways. But exploring this would force us to wade into the muddy waters of the relationship between reasons and oughts, and is beyond the scope of this project.

In defense of Perspectivalism, Lord (2015, p. 35) offers the following argument, a version of the attractive strategy.

Lord's Perspectivalism Argument¹⁹

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* possesses the right reasons."

Lord's first premise is a version of OIC—it posits an ability-constraint on obligation. And it's plausible on its face. Though RRAC is not as obviously uncontroversial as (a more generic) OIC appears to be, I suspect that most metaethicists would be inclined to endorse something like it.²⁰ Even the Objectivist—Lord's chief opponent—can accept RRAC given an expansive enough understanding of ability.²¹ But despite RRAC's *prima facie* plausibility, Lord does not simply help himself to it; he mounts 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."²² The thought here is that if RRAC is false, then an agent might be obligated to do something that she could do only unintentionally, by mistake, or for some reasons other than those which genuinely call for the action. But, Lord insists, this isn't how obligation works—it can never be the case that the *only* way to do what you ought is to "get lucky." 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

¹⁹ "Right Reasons Ability Condition" is Lord's label. "Possession" is my label. Again, 'possesses' is just shorthand for 'bears the right epistemic relation to'.

²⁰ Though 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.

²¹ Lord (2015, p. 38) 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." As we'll see, though, this is the very thing that causes problems for his strategy and others like it.

²² The falsity of RRAC has this implication given the following principle:

"Sensitive No Accident: $A \phi$ s for the right reasons just in case $A \phi$ s non-accidentally because *A* is sensitive to the right reasons" (Lord 2015, p. 36).

The relevance of Sensitive No Accident to Lord's argument will be discussed more thoroughly at the end of the next section.

one ought to ϕ even though one is unable to ϕ in a way that would be creditworthy." Why? Because one deserves credit for doing what she ought only if she does what she ought for the reason(s) she ought.²³ And if an agent can be obligated to do something despite not being able 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 possibly do in a creditworthy way. But, Lord maintains, it is *always* possible to do what one ought in a creditworthy way—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 kind of) epistemic contact with those reasons. In defense of this, Lord (2015, pp. 38-40) 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."

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, Lord argues, if we insist (against intuition) that D-Andy acts for the right reasons, then we're committed to thinking (i) 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); (ii) that his act is creditworthy (because he does what he ought for the reason that he ought); and (iii) that his action is justified (because doing what one ought for the reason that one ought explains why that action is justified).²⁴ But none of this seems

²³ This is half of Lord's Credit:

"**Credit:** A's ϕ -ing is creditworthy just in case A ϕ s for reasons that make ϕ -ing permissible" (Lord 2015, p. 37). Credit will also be discussed more thoroughly at the end of the next section.

²⁴ (i) follows from Sensitive No Accident (see footnote 22). (ii) follows from Credit (see footnote 23). (iii) follows from the Explanatory Condition:

right—D-Andy isn't sensitive to the reasons that bear on his situation, he doesn't deserve credit, 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 the Andys is that S-Andy *knows* his wife is cheating on him while D-Andy does not. 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 reason. But D-Andy doesn't have access to the fact that is the right reason for him to get a divorce. Of course 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.²⁵ Because there is no sense in which D-Andy grasps the fact that is the right reason for him to get a divorce, he does not *act for* the right reason.

From all of this, Lord concludes that the best explanation of why only S-Andy acts for the right reason is that only S-Andy bears the right epistemic relation to the right reason. Moreover, Lord continues, not only does D-Andy not *actually* act for the right reasons, he doesn't seem *able* to do so—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. Lord argues for the existence of a subjective (epistemic) constraint on obligation from a version of OIC (RRAC) and a substantive claim about the ability to act for the right reasons

“Explanatory Condition: If A ϕ s for a normative reason r, r provides a justificatory explanation of why A ϕ s” (Lord 2015, p. 34).

While the Explanatory Condition may be controversial, it is beyond the scope of this paper to critically examine it.

²⁵ Epistemically lucky, of course. The fact that his wife is cheating on him is decidedly unlucky.

(Possession, an instance of SCA). Again, Lord isn't the only one who pursues this kind of strategy. As mentioned in the previous section, a host of philosophers have offered arguments along these lines. Williams' (1981) famous argument for Reasons Internalism can be read this way—as arguing that R is a reason for an agent to φ only if that agent is able to φ for R (OIC), and that an agent is able to φ for R only if there exists a sound deliberative route from her subjective motivational set, through R, to φ -ing (SCA).²⁶ Korsgaard (1986) and Joyce (2001, ch. 5) take a similar approach along the way to arguing, respectively, against skepticism about practical reasoning and for moral error theory. Darwall (1992) employs the attractive strategy in defense of Constitutivism—arguing that a person ought to φ only if she is able to freely choose to φ (OIC), and that an agent is able to freely choose to φ only if she can be motivated to φ through practical reasoning (SCA).²⁷ Huemer (2005, ch. 7) concedes the attractive strategy in his defense of ethical intuitionism—admitting that an agent is obligated to φ only if she is able to φ (OIC) and an agent is able to φ only if she can be motivated to φ (SCA).²⁸

So I take Lord's argument for Perspectivalism to be just one (especially clear and explicit) instance of a common line of reasoning. But despite the appeal and popularity of the attractive strategy, I worry that it equivocates on multiple senses of ability. It's the task of the next section to explain why.

²⁶ Anomaly (2008) explicitly argues that Williams should be read this way. It's also worth mentioning that Williams' argument is in terms of reasons rather than obligations. But similar considerations apply.

²⁷ Heathwood (2011) urges this reading of Darwall.

²⁸ Notably, Huemer goes on to argue *against* the existence of a (substantive) subjective constraint on obligation on the grounds that agents can be motivated by beliefs about impartial reasons.

IV. WHAT MAKES AN AGENT ABLE?

I admit to finding the attractive strategy attractive. But conspicuously absent from implementations of it is any serious discussion of what *ability* amounts to.²⁹ And this is worrisome. While each premise looks plausible, and is often well defended, there is no consensus on the right way to analyze ability. This worry is compounded by the fact that the truth of an ability ascription is, in at least one important respect, (almost always) context sensitive. Given this, it's important to ensure that the sense of ability operative in OIC is the same sense of ability operating in the accompanying SCA in any instance of the attractive strategy. But once we appreciate the variability of ability ascriptions, it becomes clear that different metaethical accounts take different senses of ability to constrain obligation. And there's typically little or no case made for thinking that the sense of ability appealed to in OIC is the same as the sense of ability operative in SCA. The worry, in other words, is that given the context sensitivity of ability ascriptions, we need some reason to think that the same sense of ability is operative in each premise of the attractive strategy, and typically no such reason is offered. In light of this, the opponent of the attractive strategy's conclusion can accept that *some sense* of ability constrains obligation, and *some sense* of ability is constrained by an agent's subjective perspective, all while insisting that that these are simply *different* senses of ability.

As a first step toward demonstrating the context sensitivity of ability ascriptions, consider 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.

²⁹ To his credit, 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 (Lord 2015, p. 34, fn 11; p. 40, fn 20). But I take the worries expressed here to apply despite these clarifications.

All three of these claims might be true—the canyon might be too wide for Katie, Clara might be too embarrassed to face her mom, and Luke might be already committed to babysitting his niece. But, crucially, what *makes* each claim true is different in each case. In Ab_1 , Katie's inability would have to be explained by something like the fact that there is no possible world in which she (i) is in (more or less) identical external circumstances, (ii) has (more or less) identical physical traits and talents, and (iii) actually leaps the canyon.³⁰ (Conversely, her *ability*, if she had it, would be explained by the *existence* of such a world.) In Ab_2 , Clara's inability couldn't be similarly explained—clearly there exists a world in which she is in (more or less) identical external circumstances, with (more or less) identical physical traits, and looks her mother in the eye. Instead, Clara's inability would have to 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 Ab_3 , Luke's inability demands a different explanation still—his would have to 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 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 (2012) convincingly argues that ability claims are always relational in this way—they relate some event (leaping the canyon, looking someone in the eye, attending the party) to some (implicitly or explicitly) restricted set of worlds.³¹ 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 agent and event. But the restriction is often implicit (as it was in Ab_{1-3} , and almost always is in natural

³⁰ When I speak of an agent's physical traits, I mean *merely* physical—i.e. physical, but not mental, epistemic, or normative (this might require a rather course-grained understanding of physical traits). Nothing I say here turns on whether or not epistemic or normative properties supervene on physical properties.

³¹ Kratzer's discussion is in terms of "must" and "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 purposes here.

language), and when it's implicit, it must be derived from context. Usually this is no trouble at all. For example, Ab_{1-3} might very naturally be understood as implicitly expressing the following *given* clauses:³²

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.³³ 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.

Before moving on, I should note that, for ease of explication, I've been fudging one of Kratzer's more important insights. It is 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' (and its variants) 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 the meaning of 'ability' is static. It's the *relata* of ability ascriptions that are variable. With this caveat in mind—again, for ease of explication—I'll continue to speak in terms of different "senses" of ability.

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 each premise of the attractive strategy. Again, the worry is that the kinds of considerations that

³² Kratzer prefers "in view of" rather than "given". But this is just a stylistic preference.

³³ 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.

support OIC don't obviously require a sense of ability that is constrained by an agent's subjective perspective. The point is a bit difficult to see in the abstract. It's best made, I think, demonstratively—by showing how the worry arises for a particular instance of the attractive strategy. Hence the focus on Lord. The task now is to determine what sense of ability is operative in RRAC and Possession.³⁴

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 to act for the right reason, *given* what he 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 external 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 epistemic profile. So the implicit restrictor in Possession limits the possibility space to externally and *epistemically* similar worlds.³⁵ 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 traits and external and epistemic circumstances are (more or less) the same, and that agent intentionally ϕ s.³⁶

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 and talents. Also importantly, this isn't a

³⁴ In case you forgot:

Right Reasons Ability Condition (RRAC): "If A ought to ϕ , then A has the ability to ϕ for the right reasons."

Possession: "If A has the ability to ϕ for the right reasons, then A possesses the right reasons."

³⁵ I mean for external and epistemic similarity to come apart. So if you're an externalist about epistemic justification, you should read "externally similar" as sufficiently coarse-grained to allow for epistemic variation.

³⁶ The standard view seems to be that obligation entails the ability to do things *intentionally* (whatever sense of ability we have in mind). 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 here (though be aware that it will do some work in section V).

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. But she can do all this while insisting that it's just false 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 traits and external 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 traits and external 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.

I want to emphasize here that not only can the Objectivist and Constructivist accept *both* RRAC and Possession (while insisting that a different sense of ability is operative in each), they can

do so *for the very same reasons* that Lord does. This point is worth dwelling on for a moment because it illustrates just how deeply the problem of the context sensitivity of ability ascriptions can infect the attractive strategy.³⁷

Lord (2015, pp. 35-38) supports RRAC by pointing out that its negation has uncomfortable implications given the following two principles:

Sensitive No Accident (SNA): “A ϕ s for the right reasons just in case A ϕ s non-accidentally because A is sensitive to the right reasons.”

Credit: “A’s ϕ -ing is creditworthy just in case A ϕ s for reasons that make ϕ -ing permissible.”

Specifically, if RRAC is false, then (given SNA) (i) it is possible for an agent to be obligated to ϕ even though she is unable to ϕ non-accidentally in the right way, and (given Credit) (ii) it is possible for an agent to be obligated to ϕ even though she is unable to ϕ in a way that would be creditworthy. Since (i) and (ii) are implausible, RRAC must be true.³⁸ The problem for Lord is that neither SNA nor Credit nor the denial of (i) and (ii) requires a particular reading of ‘ability’.³⁹ Given this, the Objectivist and Constructivist can comfortably accept both SNA and Credit, comfortably deny (i) and (ii), and thereby endorse RRAC, while still reading in their preferred sense of ability.⁴⁰ So Lord’s arguments for RRAC cannot alone establish the Perspectivalist’s required sense of ability.

³⁷ I’m indebted here and in the remainder of this section to an anonymous referee who pressed me to clarify this point.

³⁸ Again, any of these moves might be questioned (though I admit to finding them all quite plausible). But since the current task is to show that Lord’s opponents *needn’t* question them to deny his conclusion, I set any other criticisms aside.

³⁹ To be clear, since (i) and (ii) are ability claims, denying them requires *some* understanding of the operative sense of ability operating. The point here just is that denying (i) and (ii) doesn’t favor any of the candidate senses of ability under consideration (O-, C-, and P-) over any other.

⁴⁰ Notably, this would commit the Objectivist and Constructivist to thinking that being obligated to ϕ entails the O- and C-ability (respectively) to be sensitive to the right reasons and to ϕ in a way that is creditworthy, given the following very plausible principle:

Ability Substitution: For any sense of ability, if ϕ -ing is necessary and sufficient for ψ -ing, then an agent is able to ϕ if and only if she is able to ψ .

But these implications are unobjectionable, and there’s no reason for either the Objectivist or Constructivist to be unhappy with them.

But perhaps they were never meant to do so. Instead, Lord might insist that the reason to think that the sense of ability in RRAC is epistemically restricted comes from thinking about RRAC *in conjunction* with his arguments for Possession. Specifically, since his arguments for Possession appeal to the same principles that support RRAC, we might read Lord as implicitly putting pressure on his opponents to think that the sense of ability operative in RRAC must be the same as the sense operating in Possession (which, again, we should all agree is P-ability). But, upon further review, the Objectivist and Constructivist needn't be bothered here, either.

Recall that Lord's support for Possession runs through his consideration of the Andys. Again, that S-Andy acts for the right reason (when divorcing his wife), while D-Andy does not, is intuitive. Importantly, though, this judgment is also supported by SNA and Credit.⁴¹ The thought is that if D-Andy does act for the right reason, then (given SNA) (i) D-Andy is sensitive to the right reason, and (given Credit) (ii) D-Andy's action is creditworthy. But neither of these implications is plausible. So D-Andy does not act for the right reason.

With all this, the Objectivist and Constructivist can happily agree. They can agree that D-Andy doesn't act for the right reason. And, since they can accept SNA and Credit, they can also agree that this is so precisely because *actually acting* for the right reason entails that the acting agent is sensitive to her reasons and that her action is creditworthy, neither of which is true of D-Andy. Since, as Lord (2015, p. 39) points out, "the only relevant difference between the two Andys is that ... the relevant fact is within [S-]Andy's epistemic ken," the Objectivist and Constructivist can even agree that an agent (actually) acts for the right reason only if that reason is in her epistemic ken.⁴² So up to this point, there's nothing for the Objectivist and Constructivist to even quibble with.

⁴¹ Lord also appeals to the Explanatory Condition in support of the claim that D-Andy doesn't act for the right reason (see footnote 24, above) (Lord 2015, p. 34). But since the explanatory condition doesn't play a direct role in Lord's defense of RRAC, it's not relevant to the response currently under consideration.

⁴² Note that since there is not yet a reference to ability, this is not yet Possession.

Lord's next move is to insist that beyond not *actually* acting for the right reasons, D-Andy isn't *able* to act for the right reasons. But here again, no one should disagree. There is certainly a sense of ability according to which D-Andy is unable to act for the right reasons—specifically, a sense of ability that takes D-Andy's epistemic state to be part of the (implicit) restrictor. The Objectivist and Constructivist can accept all this while shrugging and insisting (perhaps under their breath), "But of course given the *normatively relevant* sense of ability [O- or C-], D-Andy *is* able to act for the right reasons, and (given SNA and Credit) thereby *is* able to be reasons-sensitive and to act in a way that is creditworthy."⁴³ The Objectivist and Constructivist, then, can maintain that there's a (perfectly legitimate) sense in which D-Andy is *unable* to act for the right reasons, and a (normatively relevant) sense in which D-Andy is *able* to act for the right reasons. And they can do this without denying any of the principles that motivate Lord's case or saddling themselves with any obviously awkward implications.⁴⁴

⁴³ Perhaps this thought needs some more fleshing out. So here it is in a bit more detail. Our Objectivist and Constructivist each accept RRAC and insist that the operative sense of ability in it is O- and C-ability, respectively. Since each *also* accepts SNA and Credit, and since SNA and Credit express necessary and sufficient conditions for acting for the right reason, they are also committed to thinking that the ability to act for the right reason entails the abilities to be sensitive to the right reasons and to act in a way that is creditworthy (given Ability Substitution, [see footnote 40](#)). D-Andy might initially appear to be a counterexample to these commitments. But he's not. Because just as it's uncontroversially true that D-Andy is unable to be sensitive to the right reasons and to act in a creditworthy way on *some* senses of ability, it's uncontroversially false on *other* senses of ability. The Objectivist and Constructivist just insist that the normatively relevant sense of ability is one of those other senses. So, the Objectivist and Constructivist follow Lord in maintaining that D-Andy is unable to act in a reasons-sensitive and creditworthy way, but only by inserting a sense of ability that each takes to be irrelevant to RRAC. And they can do this without incurring any obvious cost.

⁴⁴ There is one additional move Lord might make. At the end of his discussion of the Andys, in support of his claim that D-Andy isn't able to act for the right reasons, Lord mentions that D-Andy "seems to be exercising all the abilities he has" and "isn't ... holding anything back" (Lord 2015, p. 39). There are two things to be said about this thought. First, and familiarly, the Objectivist and Constructivist might continue to shrug—insisting that this intuitive sense in which D-Andy isn't holding anything back in this case just isn't tracking the normatively relevant sense of ability. But, more speculatively, we might also read Lord here as gesturing toward some *criteria* for determining the normatively relevant sense of ability—that it *should* track our intuitive grasp of when an agent is or isn't "holding back." This passage is suggestive. To the extent that this would provide us with some reason for thinking that a particular sense of ability is the normatively relevant one, we might view Lord as taking a preliminary step toward making the kind of argument endorsed here in this paper. But, of course, much more would have to be said. (Thanks to the anonymous referee for pointing out this strategy.)

The lesson of all this is that the problems caused for the attractive strategy by the context sensitivity of ability ascriptions are not merely surface level—it can taint the premises as well as their justification. In Lord's case, his opponents can accept both RRAC and Possession for all of the reasons he endorses them, all while reasonably insisting that a different sense of ability is operating in each premise. The point here isn't that there's any problem with Lord's view—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.

This problem for Lord is a problem for the attractive strategy generally. The kinds of considerations cited in support of (versions of) OIC—when it's not simply assumed—typically don't require a subjectively restricted sense of ability. And this isn't surprising. To do so would be to beg the question against those who deny that obligation is subjectively constrained or insist that it is subject to some competing constraint. But the attractive strategy continues—via SCA—by insisting that the relevant ability is *in general* subjectively constrained (in some way or another). And this can be true only given a *subjectively restricted* sense of ability—a sense of ability, in other words, that is dialectically ineligible to appear in OIC without some further argument that that particular sense of ability constrains obligation. So the problem isn't unique to Lord's argument for Perspectivalism. It applies to any attempt to deploy the attractive strategy that doesn't explicitly grapple with the sense of ability that operates in the relevant version of OIC.

So, despite first appearances, the attractive strategy (on its own, anyway) doesn't seem very well-suited to advance the metaethical debate. While most metaethicists are happy to endorse some version of OIC, it doesn't represent the point of consensus that it seems to. Different metaethical approaches take different and *inconsistent* senses of ability to be operative in OIC. Because of this, appeals to substantive claims about 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.

V. SO WHAT CAN ABILITY DO?

Does all this mean that thinking about ability is useless in metaethical contexts? 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 a metaethical debate. 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, in my view, is not to abandon ability talk altogether, but to offer a defense of a particular sense of ability *as the normatively relevant one*. In this section, I make a very 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 circumstances beyond her control.

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 answer the central deliberative question, '*What should I do?*'. If a sense of ability characterizes an agent as able-to-perform some action that *cannot* answer the central deliberative question—in the sense that the central deliberative question remains open despite the answer—then we've good reason to think that that sense of ability isn't restrictive enough.

To see this strategy in action, first take a sense of ability that almost everyone would agree is *not* restrictive enough.

G-ability: An agent is G-able to φ just in case there is some possible world in which the agent's physical traits are (more or less) the same, and that agent intentionally φ s.⁴⁵

G-ability, like O-, P-, and C-ability, is a perfectly respectable sense of ability. It captures the sense in which an agent might be able to do something *given her traits and talents*, even if she currently lacks the *opportunity* to do so. For example, consider the following ability ascription:

Ab₄ – Doug is able to play the piano.

Ab₄ may be true even if Doug is currently driving his granddaughter to Cleveland. It's true in virtue of the existence of some possible world in which Doug, with (more or less) the same physical traits and talents, intentionally plays the piano. Moreover, it's true despite the fact that there is *no* world in which Doug plays the piano *given* his current external circumstances, since his current external circumstances are not part of the (implicit) restriction.

Despite being a respectable sense of ability, G-ability is clearly *not* up to the task of capturing the full extent to which abilities constrain obligations. To see why, imagine that, while driving, Doug's granddaughter starts to cry. Suppose further that the sound of piano music would soothe her. We could tell Doug that he should play the piano to calm her down. But clearly this advice leaves something to be desired. To it, he might reasonably reply: "I see how that would be nice. But since I can't do that right now, *what should I do?*" The fact that Doug can still ask the central deliberative question suggests that the action we offered failed to answer it. So playing the piano can't be the action which Doug is obligated to perform. Moreover, Doug's explanation of *why* our recommendation was bad seems right— "play the piano" is not a viable answer to the central deliberative question because Doug isn't *able* (in the relevant sense) to do so. There are features of Doug's circumstances, beyond his control (most relevantly, the lack of a piano in his car), that prevent him from intentionally performing the recommended action.

⁴⁵ I call this "G-ability" as a nod to Mele's (2003) characterization of "general ability". His (insightful) project is primarily concerned with analyzing ability, not with exploring its normative implications.

“Play the piano” is not a viable answer to the central deliberative question because, given Doug’s circumstances, doing so isn’t an option. But, again, Doug *is* G-able to play the piano. So, G-ability is not sufficiently restricted to capture the extent to which our abilities limit our obligations.

Similar reasoning applies, I think, to O-, C-, and P-ability. Take 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 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?⁴⁶

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?⁴⁷

Let’s work with Miners first. We might recommend to Billy that she 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?*” As in Doug’s case, we’ve again failed to answer the central deliberative question. And, again, the explanation seems to be that we’ve cited an action which Billy is not *able* (in the relevant sense) to perform. There are features of Billy’s circumstances, beyond her control (most relevantly, the unavailability of evidence about the miners’ location), that prevent her from intentionally performing the recommended

⁴⁶ This is Lord’s (2015) statement of the case. It was conceived by Donald Regan (1980) and made famous by Derek Parfit (2011). See Kolodny and MacFarlane (2010) for a recent, influential discussion of the case.

⁴⁷ This is Kieran Setiya’s (2007) statement of the case (with some slight, non-substantive modifications). It was conceived by Gary Watson (1975).

action.⁴⁸ “Sandbag the shaft the miners are in” is not a viable answer to the central deliberative question because, given Billy’s circumstances, doing so just isn’t an option. But clearly Billy *is* O-able and C-able to sandbag the shaft that the miners are in—there exists some possible world with (more or less) identical external 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.

Is P-ability enough? It seems to fare well in Miners.⁴⁹ It correctly categorizes Billy as unable to sandbag the shaft the miners are in—she isn’t P-able to sandbag the shaft that the miners are in because there’s no possible world in which she is in (more or less) the same epistemic circumstances and intentionally does so. 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, very plausibly, the explanation for this failure is that we’ve cited an action which Stan is not *able* (in the relevant sense) to perform. There are features of Stan’s circumstances, beyond his control (most relevantly, his overwhelming desire to knock in his rival’s teeth), that prevent him from intentionally performing the recommended action. “Shake your opponent’s hand” is not a viable answer to the central deliberative question because, given Stan’s circumstances, doing so just isn’t an option. But clearly Stan *is* P-able (and O-able, incidentally) to shake his opponent’s hand—there exists some possible world with (more or less) identical external and epistemic circumstances in which Stan shakes his

⁴⁸ Note here, that the intentionality of the action is doing some important work. To intentionally sandbag the shaft the miners are in Billy would (at least) need a justified belief about which shaft they were in. So although the features of Billy’s circumstances (beyond her control) don’t prevent her from blocking the shaft the miners are in *by accident*, they do prevent her from doing so intentionally. But none of this strikes me as problematic. As mentioned in footnote 36, I’m inclined to think that obligations are obligations to act intentionally. But (again) defending this is beyond the scope of this paper.

⁴⁹ This isn’t surprising. Miners is one of the cases Lord uses to motivate Perspectivalism.

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 our abilities on some sense of ability that is *both* epistemically *and* motivationally restricted. Only such a dually restricted sense of ability could get *both* Miners and Squash right. Precisely how we should characterize each of these constraints is beyond my purposes here. My more modest goal is just to show that both *kinds* of restrictions are required. 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. And, as Miners and Squash highlight, an agent often *lacks* control over epistemic and motivational features of her circumstances that are relevant to the central deliberative question, '*What should I do?*'. In this sense, in order for an agent's obligations to respect the boundary between the agent and her circumstances, they must not call for her to stray too far from her epistemic and motivational lot. The central deliberative question has not been answered until an option is cited that respects the extent to which the agent is slave to (at least certain of) her attitudes.

VI. CONCLUSION

I've argued that appeals to ability obscure metaethical disagreement by concealing the commitments of different metaethical theories about the sense of ability that constrains obligation. I attempted to demonstrate this by considering an attractive strategy for defending a subjective constraint on obligation and showing that it subtly equivocates between competing senses of ability. I then endeavored to show that there might be good reason to think that a particular sense of ability is operative in (something like) OIC. Specifically, I argued that the operative sense of ability must be one that restricts the relevant possibility space to epistemically and motivationally similar worlds (determining the precise character of this

similarity being a project for another time). It's worth noting that despite endorsing a (relatively strong) subjective constraint on agents' obligations, nothing I've said implies that a complete account of agents' obligations will tell us the whole evaluative story. Even once we've answered the question of what an agent is obligated (in the sense that I'm interested in) to do, it may still be instructive to ask what would be best—or most virtuous or most rational, etc.—for her to do. And, as Doug and Billy and Stan suggest, these are different questions, and may demand different answers.

More generally, I hope to have identified a promising way to argue for a particular sense of ability as the normatively relevant one. To the extent that we can get an independent grip on the reach of agency—on the extent to which an agent exercises genuine control over her circumstances—we can appeal to such considerations to home 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.
- Darwall, Stephen L. (1992). Internalism and Agency. *Philosophical Perspectives* 6: 155-174.
- Gert, Bernard (1998). *Morality: Its Nature and Justification*. Oxford University Press.
- Graham, Peter A. (2010). In Defense of Objectivism about Moral Obligation. *Ethics* 121: 88–115.
- Graham, Peter A. (2011). 'Ought' and Ability. *Philosophical Review* 120 (3):337-382.
- 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.
- Huemer, Michael (2005). *Ethical Intuitionism*. Palgrave Macmillan.
- Joyce, Richard (2001). *The Myth of Morality*. Cambridge University Press.
- Kolodny, Niko & MacFarlane, John (2010). Ifs and Oughts. *Journal of Philosophy* 107 (3): 115-143.
- Korsgaard, Christine (1986). Skepticism about Practical Reason. *Journal of Philosophy* 83 (1):5-25.
- Kratzer, Angelika (2012). What 'Must' and 'Can' Must and Can Mean. In *Modals and Conditionals*. Oxford University Press: 4-20.
- 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.
- Mele, Alfred R. (2003). Agents' abilities. *Noûs* 37 (3):447–470.
- Moore, G. E. (1912). *Ethics*. Oxford University Press.
- Parfit, Derek. 2011. *On What Matters*. Oxford: Oxford University Press.
- 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.
- Streiffer, Robert (2003). *Moral Relativism and Reasons for Action*. Routledge.
- Streumer, Bart (forthcoming). Reasons and Ability. In Daniel Star (ed.) *The Oxford Handbook of Reasons and Normativity*. Oxford University Press.

- Thomson, Judith Jarvis (1986). Imposing Risks. in *Rights, Restitution, and Risk: Essays in Moral Theory*. Harvard University Press: 173-91.
- 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 (2009). The "Good" and the "Right" Revisited. *Philosophical Perspectives* 23 (1):499-519.
- 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.