

## **The Radical Behavioral Challenge and Wide-Scope Obligations in Business**

Hasko von Kriegstein – Ryerson University, Ted Rogers School of Management

### **Abstract**

This paper responds to the *Radical Behavioral Challenge* (RBC) to normative business ethics. According to RBC, recent research on *bounded ethicality* shows that it is psychologically impossible for people to follow the prescriptions of normative business ethics. Thus, said prescriptions run afoul of the principle that nobody has an obligation to do something that they cannot do. I show that the only explicit response to this challenge in the business ethics literature (due to Kim et al.) is flawed because it limits normative business ethics to condemning practitioners' behaviour without providing usable suggestions for how to do better. I argue that a more satisfying response is to, first, recognize that most obligations in business are wide-scope which, second, implies that there are multiple ways of fulfilling them. This provides a solid theoretical grounding for the increasingly popular view that we have obligations to erect institutional safeguards when bounded ethicality is likely to interfere with our ability to do what is right. I conclude with examples of such safeguards and some advice on how to use the research findings on bounded ethicality in designing ethical business organizations.

**Keywords:** Bounded ethicality; Ought implies can; Wide-scope obligations; Business ethics denial; Skepticism about business ethics

**The author declares that they have no conflict of interest.**

## Introduction

The field of business ethics has always been, and continues to be, subject to skeptical challenges. People who believe that there is a deep tension between ethics and business often present themselves as “realists” and accuse those who work in normative business ethics of a naïve kind of idealism (von Kriegstein 2019). This can serve as a useful reminder to take seriously the special challenges of thinking about ethics in a competitive domain like business. Normative business ethics, especially when done by philosophers, is always at risk of insisting on high-minded ideals that cannot be easily applied by practitioners. But for normative business ethics to be a useful endeavour it has to be able to provide guidance – not just tools for condemnation (though we should not be shy to condemn, when appropriate, either).

In this paper, I address a relatively new skeptical challenge to normative business ethics. Following Kim, Monge, and Strudler, I will refer to it as the *Radical Behavioral Challenge* (RBC). According to RBC, much of what normative business ethicists have to say about our obligations in a business context is misguided because it relies on unrealistic assumptions about the way in which actual people make decisions. In particular, people often make decisions based on subconscious heuristics and inclinations, and this can lead to actions that go against their own professed, and sincerely held, moral beliefs. This phenomenon is known as ‘bounded ethicality’. Insofar as people are incapable of gaining control of this subconscious decision-making apparatus, they might be incapable of following moral rules even when they fully endorse them. As Kim et al. argue, many prescriptions of normative business ethics are such that ethically bounded agents will have difficulty following them (Kim et al. 2015, 346). Thus, so the challenge, said prescriptions run afoul of the principle that ‘ought implies can’, i.e. that nobody has an obligation to do something that they cannot do.

To counter that challenge, Kim et al. propose that ‘can’ in ‘ought implies can’ refers to the abilities of a hypothetical ideally rational agent, rather than to the occurrent abilities of actual businesspeople.<sup>1</sup> But this response will not do. The proposed version of ‘ought implies can’ fails to rule out a number of obligations that clearly ought to be ruled out by the principle. Moreover, this version of the principle discards one of its most attractive features, namely its function of ensuring that obligations correspond to action-guiding reasons. I propose that, instead of tinkering with ‘ought implies can’, the best answer to RBC is to take a closer look at the form that our obligations in the business context typically take. Appealing to a distinction familiar from debates about the nature of rationality, I will argue that most dictates of normative business ethics are wide-scope, rather narrow-scope, obligations. In other words, agents are not simply under an obligation to perform this or that action, but rather they are under such an obligation ‘provided they take on certain roles or responsibilities’. This simple shift of perspective allows us to see that these obligations can be met in more than one way. Most obviously, agents who are really incapable of acting in the required ways, can avoid violating their obligations by not taking on the corresponding roles and responsibilities. More interestingly, agents may be able to manipulate themselves or their institutional environment in ways that decrease the likelihood that bounded ethicality will interfere with their ability to act how they ought to. This, too, is a way to fulfill obligations that ethically bounded agents might at first appear to be incapable of meeting. From this perspective the results of behavioural ethics do not primarily appear as fueling a skeptical

---

<sup>1</sup> As an anonymous reviewer points out, this might sound like Kim et al. are proposing something akin to the ‘reasonable person’ standard as commonly employed in law. It is important to clarify that this is not so. The hypothetical ‘rational agent’ that Kim et al. have in mind is a being very different from any actual person in that they do not suffer from any of the psychological limitations that afflict humans in general. By contrast, the ‘reasonable person’ employed in law is an ordinary human being equipped with ordinary common sense.

challenge to normative business ethics, but rather as helpful tools for further normative theorizing. Behavioural ethics does not stop at identifying and cataloguing psychological tendencies that may causally contribute to ethical lapses; researchers also explore the efficacy of various strategies to combat such tendencies or insulate our decision making from them. What they find is important, as ethically bounded agents who wish to take on certain roles and responsibilities might have an obligation to put those strategies in place.

The paper proceeds as follows. I begin with a brief sketch of the phenomenon of bounded ethicality as studied by behavioral ethics researchers (section 1), followed by some caveats and boundary-setting with regards to my use of the empirical findings in that area (section 2). In section 3, I introduce the principle that ‘ought implies can’, and show how its contrapositive, combined with the findings of behavioural ethics, gives rise to RBC. In section 4, I discuss Kim et al.’s proposed response to RBC and argue that it has serious problems. After that, I introduce the notion of wide-scope obligations and show how conceptualizing the dictates of normative business ethics as such offers a way to defuse RBC (section 5). Before concluding, I draw out some of the theoretical and practical implications of adopting this framework (sections 6 and 7).

## **1. Bounded Ethicality**

Human decision making deviates in systematic ways from the prescriptions of logic, rational choice theory, and Bayesian probability calculations. This is one of the upshots of the *heuristics and biases* research program the creation of which is typically credited to Amos Tversky and Daniel Kahneman (Kahneman and Klein 2009). Another important upshot of this research program is that much of our decision making involves unconscious processes. We often make up our minds without ever consciously considering reasons for or against a course of action. This fact is partly

obscured by our ability to rationalize. Even when our decisions are more or less completely determined by our subconscious machinery of heuristics and biases, we are usually able, after the fact, to tell a story for why we acted as we did. And those stories, produced, as they are, by our conscious mind, will very much be stories of conscious, rational deliberation.

Once we accept that conscious beliefs and attitudes have a limited effect on how we behave, we may worry that we might often behave in ways that are inconsistent with the ethical beliefs we explicitly endorse. This worry has spurred an offshoot of the heuristics and biases program called *behavioral ethics*. According to this research, people's behaviour is, indeed, often inconsistent with their conscious moral beliefs, because they are influenced by subconscious psychological biases. This phenomenon is commonly referred to as 'bounded ethicality' (Chugh et al. 2005). At a high level of abstraction, bounded ethicality consists in our subconscious biases and heuristics being imperfectly aligned with our conscious moral beliefs. This puts us at risk of unwittingly acting contrary to the latter, whenever we make decisions guided by the former.<sup>2</sup> Perhaps the most well-known class of phenomena in this category comprises various forms of egocentrism. When members of a group are asked to estimate the percentage of their contribution to the group's effort or success, for example, the sum of the individual estimates will typically be well above 100%.

---

<sup>2</sup> The term 'bounded ethicality' is an obvious reference to Simon's concept of 'bounded rationality' (Simon 1957). And indeed, there is significant overlap between these concepts and many phenomena can be classified as either bounded rationality or bounded ethicality. The main difference between the two concepts is that bounded rationality draws a contrast to the ideally rational 'homo economicus' who uses their flawless rational capacities in pursuit of self-interested goals, while bounded ethicality points to the fact that the same computational and psychological limitations also apply when pursuing altruistic goals, or attempting to follow moral principles.

This indicates that members are systematically overestimating their own contributions (Messick and Sentis 1979). Thus, when people claim, what they genuinely perceive to be, their fair share of credit, they will make demands that are out of line with what they in fact deserve (Caruso et al. 2006).

Let me give two examples to illustrate how bounded ethicality may wreak havoc in the business context. First, consider motivated blindness in an accounting context: research suggests that simply being assigned different roles (such as ‘buyer’s auditor’ or ‘seller’s auditor’) leads subjects (including professional accountants) to reach very different interpretations of identical sets of information about a company, namely ones that favour their perceived in-group. What is interesting about this result is that it appears to hold, even when subjects are strongly incentivized to be objective (Moore et al. 2010). Thus, their favouring of an (imaginary) client’s interest cannot be explained by appealing to the subject’s self-interest. It appears rather that subjects simply perceive the information differently depending on what role they are assigned, and are thus genuinely incapable of providing an objective assessment despite their best efforts. This suggests that the exhortation to evaluate business documents without favouring one’s perceived in-group might very well be beyond what ordinary people are capable of doing.<sup>3</sup> In the real world, this issue is often going to be compounded by self-interest pulling in the same direction. An accountant who perceives a company they are auditing as a potential employer, for example, might be incapable

---

<sup>3</sup> As the editor pointed out to me, some readers may think of auditor objectivity as a merely legal, rather than moral, obligation. Since accounting is a self-regulating profession, it might then be up to professional bodies to decide how strictly to interpret the requirement to be objective. However, I would argue that the requirement to be objective in one’s audits is not a free-standing legal or self-regulatory requirement. Rather, it is a moral requirement and ensuring it is abided by is part of the mandate of the accounting profession. Were the accounting profession to reject this, society would have much less reason to grant it the privileges that come with being a recognized profession (e.g. monopoly over the service offered and the right to self-regulate) (de Bruin 2015, 186).

of performing the audit objectively, even if they have the best intentions of not letting their own self-interest interfere with their performance (Zhang et al. 2015).

A similar conclusion may be drawn when considering implicit biases in the context of decisions on hiring or loan applications. Implicit biases are associations between different ideas that colour our decision making in ways we might never become aware of. For example, many people who do not endorse racist beliefs, are nevertheless quicker to associate negative concepts with black faces and positive concepts with white faces than vice versa (Bertrand et al. 2005). As a consequence, a person conducting an interview with a person from a group that they are subconsciously biased against may evaluate that person less favourably than somebody else with identical credentials. This would instantiate epistemic injustice, i.e. the phenomenon that someone is not being taken sufficiently seriously as a possessor and communicator of knowledge just because they are a member of a group that is defined by a particular characteristic such as ethnicity or gender (Fricker 2007). There is ample evidence that such forms of epistemic injustice exist in the business context and have detrimental effects on the economic prospects of members of minority groups, women, and other disadvantaged groups (de Bruin 2019). Simply noting such outcomes does not tell us whether they are the result of implicit bias or overt racism (Feldman 2018, 211).<sup>4</sup> The important point for current purposes, however, is that we can easily see how implicit bias can replicate the results of overt racism (Saul 2013).

---

<sup>4</sup> This problem is exacerbated when epistemic injustice leads to vicious cycles. Groups that are consistently subject to epistemic injustice might increasingly see efforts to obtain and communicate knowledge as fruitless. This may lead them to reduce such efforts thereby apparently validating the prejudices that started the cycle (De Bruin 2019).

## **2. Some Caveats**

Before moving on, I should emphasize that the sketch of bounded ethicality just provided is very much simplistic. For one, many of the findings of behavioral ethics are hotly disputed (see Koehn 2020 for the suggestion that the entire research program is fundamentally misguided). For instance, recent meta-studies have cast serious doubt on whether implicit bias really does have the behavioral consequences imagined at the end of section 1 (Oswald et al. 2013; Forscher et al. 2019). Measures of implicit bias have also been shown to have questionable test-retest reliability (e.g. Gawronsky et al. 2017), which might be taken as a sign either that implicit biases are not stable personality traits, or that the instruments currently used to measure implicit bias are inadequate. Consequently, there is a lively debate about the validity of research about implicit bias (see Brownstein et al. 2020 for an overview of, and response to, such criticisms), and its implications for moral responsibility (e.g., Levy 2017; Buckwalter 2018). A further simplification in the sketch above is that I left out the fact that many of the relevant phenomena come in degrees. We might, for example, be incapable of processing information in a fully unbiased, objective manner. But that is compatible with there being different degrees of bias, and us being capable of doing quite well at keeping bias at an acceptable level (at least some of the time, and under some circumstances). Finally, behavior is usually determined by a combination of various personal and situational factors. Thus, there is every reason to think that, even if the phenomena of bounded ethicality causally influence behavior, they will rarely, if ever, be the sole cause of an action (Buckwalter 2018).

For the purposes of the next three sections I will ignore these complexities, and operate with the simplified picture of bounded ethicality sketched out in section 1. This will allow me to present the most radical version of the skeptical challenge this research program presents to

normative business ethics (section 3), as well as the most prominent response the business ethics literature has produced to this challenge (section 4). As we will see, part of what makes this response inadequate is that it threatens to render the results of behavioral ethics research, and with it any nuanced understanding of them, normatively irrelevant. After presenting my own response to the skeptical challenge (section 5), I show how it is better equipped to account for the normative import of our evolving understanding of bounded ethicality (sections 6 and 7).

### **3. ‘Ought implies Can’ and the Radical Behavioural Challenge**

Imagine you are on a hiking trip and come across a person whose arm has been trapped under a rock. The rock is too heavy for them to lift with their other arm, but you could do the job. Clearly, you are morally obligated to rescue this person. Now suppose, instead, that you arrive just a split-second after the person has fallen off a 200m cliff. Here you are under no moral obligation to rescue them. What explains the difference between the two cases is, of course, that in the second case you cannot rescue the person. This illustrates:

Ought Implies Can (OIC): Necessarily, when an agent has a moral obligation to  $\phi$ , the agent can  $\phi$ .

Note that this principle is attractive mainly for two reasons. First, OIC ensures that morality cannot make absurd demands on us, such as to fly to the rescue of a person falling off a cliff, or travelling through time to prevent the horrors of the past. Second, OIC ensures that moral obligations correspond to action-guiding reasons. If we could have obligations that are impossible to abide by, the fact that I am under an obligation to  $\phi$  would not necessarily give me an action-guiding reason to  $\phi$ , since I cannot have such a reason to do what I cannot (Vranas 2007).

OIC's contrapositive ('cannot' implies 'not the case that ought') can be used to rule out putative obligations with the following form of argument.

Premise 1: Agent A cannot  $\phi$

Premise 2: OIC's contrapositive

Conclusion: Agent A does not have an obligation to  $\phi$

RBC takes this form.<sup>5</sup> Martialing the findings of bounded ethicality, the skeptic claims that humans are psychologically incapable of following the prescriptions of normative business ethics. Combining this claim with OIC renders the conclusion that humans are not obligated to follow these prescriptions. In terms of the earlier example: it is not the case, *pace* normative business ethics, that people have an obligation to produce unbiased accounting reports, since people are psychologically incapable of doing so. The challenge is rather radical since bounded ethicality impedes a large number of tenets of normative business ethics many of which call for the kind of unbiased objectivity that we are apparently incapable of (Kim et al. 2015, 346).<sup>6</sup>

The intuitive response to RBC might be to dispute that humans are incapable of following the prescriptions of normative business ethics. That we are biased does not make it literally

---

<sup>5</sup> For a detailed investigation of various versions of this argument-form specific to implicit bias, see Holroyd 2012.

<sup>6</sup> An anonymous referee helpfully points out that there are large swaths of normative business ethics that do not seem imperiled by bounded ethicality in the way RBC suggests. The examples provided by the referee are "don't pollute, don't undermine fair competition, don't lie in your advertisements" – all surely central parts of the canon of normative business ethics. Nonetheless, I share Kim et al.'s worry that RBC, if successful, would undermine very significant parts of normative business ethics, because in business (though, of course, not only there) we are often supposed to adhere to particularly stringent standards of objectivity. Doing audits is perhaps the most obvious example. More generally we might think that the normative justification for our current economic system assumes an ideal of fair competition that is threatened when decisions are based on biases that systematically (dis)advantage some participants.

impossible to evaluate people and situations in an objective way. It may just require us to work a bit harder (see, e.g., Graham 2011). As stated, this response is too sanguine about our capacity to take control of our subconscious decision-making machinery. It is well-documented that people have a really hard time believing that they themselves are subject to all these biases (Pronin et al. 2004; Rachlinski et al. 2008), and the intuition that we can resist our biases, if we really want to, might itself be another bias (Kahneman and Klein 2009, 523). On the other hand, as pointed out in section 2, there is good reason to think that the strongest version of RBC will be unsound, because empirical findings (to date) do not back up the claim that bounded ethicality presents insurmountable barriers to ethical behavior. As Wesley Buckwalter has recently argued about implicit biases, it is much more plausible to think that these phenomena make it “more difficult but not impossible to act as morality requires” (Buckwalter 2018, 2984).

Nevertheless, I agree with Kim et al. that it is conceptually helpful to address the most radical version of RBC, and am thus granting their premises for the purposes of the next two sections. Once we have the best model in hand of how to deal with the most disconcerting empirical findings we can imagine, we can apply this model to the considerably less shocking actual findings that psychological research has produced to date.

#### **4. Tinkering with ‘Ought Implies Can’**

Even if we assume that bounded ethicality makes it literally impossible for us to follow some moral rules, we may very well wonder whether this is the kind of inability that ‘cannot’ in OIC and its contrapositive refers to. Kim et al. say that it is not. They argue that interpreting OIC as absolving agents from obligations they are incapable of meeting because of their psychological makeup would let them off the ethical hook too easily. If we accept such an interpretation of OIC, they

point out, we can no longer say that accountants ought to do objective audits, drunk people ought to drive carefully, and racist loan officers ought to look at loan applications impartially.

Following a prominent strand of interpretation of Kant's thinking on OIC (Stern 2004), Kim et al. suggest that 'can' in OIC should, instead, be interpreted as referring to the capabilities of "a rational human agent acting out of respect for the moral law and acting consistently with the laws of nature" (Kim et al. 2015, 348). This interpretation of OIC diffuses the skeptical challenge. The psychological barriers discovered by behavioural ethicists do not apply to ideal rational agents. Thus, the bounded ethicality of actual human agents does not license an invocation of OIC's contrapositive to reject the principles of normative business ethics. The question, then, is which version of OIC is more plausible: treating 'can' as referring to the occurrent ability of the actual agent as required by RBC, or as referring to the abilities of Kim et al.'s idealized agent?

To motivate their version of OIC, Kim et al. point to the extremely counterintuitive results of the alternative. This, however, is an odd move in countering a skeptical challenge. What makes RBC a skeptical position is precisely the willingness to accept the type of counterintuitive results that Kim et al. want to use as a motivation for rejecting the interpretation of OIC that leads to them. In essence, the response is to assume that the skeptic themselves would not embrace the most radical consequences of their view (e.g., that racists are exempt from impartiality obligations). This response begs the question against a skeptic who is prepared to bite the bullet on those cases. The rejection of the interpretation of OIC that licenses the inference from bounded ethicality to RBC is, thus, not well-motivated in the relevant dialectical context. In deciding whether RBC succeeds we cannot take the fact that it is counterintuitive as a tool in adjudicating between different versions of OIC. Thus, this argument for Kim et al.'s interpretation of OIC fails.

Moreover, there are positive reasons to reject Kim et al.'s interpretation of OIC. Most importantly, recall that one of the reasons OIC is an attractive principle is that obligations are supposed to correspond to action-guiding reasons: an obligation to  $\phi$  is a weighty reason to  $\phi$ . Thus, informing someone that they have an obligation to  $\phi$  is usually going to help them decide what to do. OIC is supposed to secure this role for obligations: informing someone that they have an obligation to do something they cannot, does not (straightforwardly) help them decide what to do. But under Kim et al.'s interpretation, OIC cannot play this role. Knowing that I am obliged to do something that I could do, if I was ideally rational, but in fact cannot, does not help me decide what to do. Thus, Kim et al.'s interpretation of OIC jettisons one of the most attractive features of OIC. Because of that, Kim et al.'s response to RBC amounts to the kind of naively high-minded idealism skeptics like to accuse business ethicists of. Ultimately, what Kim et al. say is: 'you are bound by the prescriptions of normative business ethics whether you are able to follow them or not!' This might allow us to satisfyingly condemn the behaviour of practitioners, but it does not provide usable advice for how to do better.

The other reason that OIC is attractive is that judging a person's conduct unethical for failing to meet a standard they could not have met seems inappropriate. The propriety of such judgements is the measuring rod that Kim et al. are employing in settling on their preferred interpretation of OIC. After rejecting the interpretation presupposed by RBC for being overly lenient, they consider whether 'can' may refer to the abilities of a rational agent acting out of respect for the moral law. Noting that this interpretation would not rule out obligations involving performing physically impossible feats, they find that it fails for the opposite reason of the first interpretation: it is too harsh (Kim et al. 2015, 348). This leads them to add the clause about acting consistently with the laws of nature. What they hope to have achieved, then, is to arrive at an

interpretation of OIC that allows us to say that biased accountants and drunk drivers are violating their ethical obligations, without committing us to obligations to, say, bring unjustly executed people back to life.

Their interpretation of OIC delivers Kim et al.'s desired results in cases involving bounded ethicality without issuing obligations that would involve breaking the laws of nature. However, it might still violate the spirit of OIC by leaving potential obligations untouched that should be ruled out. Consider the following thought experiment.

*MathBomb*: Adelaida finds a bomb, about to go off in 30 seconds, that will kill innocent people and can be defused only by punching in a code. A note on the keyboard says that the code is the square-root of  $(3323344 + 363831)/147487$ . Without a calculator ready Adelaida is unable to figure out the code in time and fails to defuse the bomb.

In this scenario, OIC should help us to say that Adelaida did not violate an obligation. But it is not clear that Kim et al.'s interpretation can deliver this result. After all, the math would be easy for an ideal rational agent. This shows that Kim et al. need to further clarify the notion of rationality employed in their version of OIC. To deal with cases like *MathBomb*, they need a conception of rationality that does not entail that a rational agent is equipped with flawless reasoning capacities in all domains (e.g. such an agent may still lack superior math skills). At the same time, to defuse RBC they need to say that a rational agent would not be subject to the nefarious influences of bounded ethicality. While such a conception of rationality could in principle be provided, doing so would not be a trivial task. To see this note, first, that it would not do to simply state that the notion of rationality employed implies only the absence of psychological tendencies leading to ethically improper behaviour. To say this would beg the question, in the current context, since the concept of rationality is supposed to help us figure out what kind of behaviour is ethically improper in the first place. Thus, one would have to specify which aspects of human rationality would be

idealized and which would be kept constant without taking cues from our intuitive judgements as to what kind of thing we are obligated to do despite our occurrent limitations.<sup>7</sup> This, however, will be difficult to do because the human limitations underlying bounded ethicality are not specific to the moral domain. In other words, bounded ethicality is not the result of challenges to being moral specifically, but rather of our limits in charting ideally rational courses of action more generally. Thus, it will be difficult to carve out a conception of a rational agent that is still subject to many of the limitations of human rationality but not affected by the specific phenomena that lead to what we intuitively think of as ethically improper behaviour.

To sum up, Kim et al.'s interpretation of OIC undermines both features that make OIC attractive. It does not preserve OIC's function of ensuring that obligations correspond to action-guiding reasons; and, pending further specification of their concept of rationality, it is unclear whether the proposal has the resources to explain why, in *MathBomb*, Adelaida has no obligation to perform a mathematical operation that is beyond her abilities. This leaves the proposal with only one strength: it allows us to save normative business ethics from the radical behavioral challenge. But in the context of an argument against RBC this consideration is not admissible. Thus, the attempt to defuse RBC by adjusting the interpretation of 'can' in OIC appears unacceptably ad hoc.

---

<sup>7</sup> To illustrate, imagine a dispute between Kim et al. and a skeptic. The skeptic claims that when they discriminate based on implicit biases they do not violate an obligation, because they cannot help but do so. Kim et al. reply that the obligation still exists because, if the skeptic was rational, he could avoid such discrimination. The skeptic charges that this reply entails that Adelaida has an obligation to defuse the *MathBomb*, since she could do that if she was rational. Kim et al. now reply that this is not the kind of rationality they have in mind. At this point the skeptic will demand to know more about the concept of rationality being used, and will insist, rightly, that the answer is justified without reference to the disputed judgement about the propriety of discrimination based on implicit biases.

## 5. Business Ethics Obligations are Wide-Scope

To apply the lesson from *MathBomb* to a case closer to our current interests, consider an accountant unable to do an accurate audit because they are bad at maths. Does this person have an obligation to do an accurate audit? yes, according to Kim et al.'s interpretation of OIC; no, according to the interpretation needed for RBC. Which is more plausible? My intuitions are unclear and I find myself wondering how this person got to be an accountant in the first place. After all, whatever we may think of the accounting profession's ability to enforce objectivity, it seems likely that it is more than capable of ensuring that nobody with inadequate maths will be certified. One possibility is that a competent accountant is, temporarily or permanently, incapacitated. Perhaps they have suffered a minor stroke, or are on vacation and quite drunk. Why would they perform an audit under such conditions? Perhaps they are somehow pressed into service (an audit needs to be done quickly and they are the only person there with the required credentials). In that case, while they might be obligated to try to pull it together and do the best they can, it seems to me that (much like Adelaida) they have no obligation to do what they cannot. By contrast, if they volunteered to do an audit despite their compromised state, the opposite seems plausible, and if they produce an audit riddled with problems stemming from math-errors, we should see them as having violated their obligation. If that is right, fine-tuning the concept of 'can' in OIC is not a promising strategy, for we can stipulate that there is no difference in terms of ability between the two cases that need to be distinguished here.

Thus, I propose to take a different route. First, note that most prescriptions of normative business ethics are hypothetical imperatives, i.e. they arise only if and when we take on certain roles or objectives. Moreover, since taking on these roles or objectives is usually not morally obligatory, it is best to think of the resulting obligations as wide-scope, rather than narrow-scope

(Kolodny 2005). Thus, the obligation in our example is not ‘do an accurate audit’, but rather ‘if you do an audit, do it accurately’. By contrast to the former, the latter can be met in two ways: by doing an accurate audit, or by not doing an audit at all. The person with inadequate math-skills cannot do the accurate audit, but, in ordinary circumstances, they can avoid doing the audit. Thus, they *can* meet their wide-scope obligation.

The same is true in cases of bounded ethicality. Our obligations are not narrow-scope (‘provide an unbiased report’, ‘look at loan applications impartially’), but rather wide-scope (‘if you are doing an audit, do an unbiased audit’, ‘if you are a loan-officer, look at loan applications impartially’). These wide-scope obligations can be fulfilled by ethically bounded agents – all they need to do is recuse themselves.<sup>8</sup> Thus, to say that human agents, though ethically bounded, are nevertheless subject to the prescriptions of normative business ethics, we do not need to weaken OIC. Instead of saying that people who are psychologically incapable of avoiding ethically problematic biases ought nevertheless do so, we can simply say that such people should avoid situations where such biases cause problems.

One might object that avoiding situations in which bounded ethicality causes problems is not as easy as I made out. An accountant who realizes that they are unable to do accurate audits may, after all, be in a position where they can scarcely afford to lose the income their position affords them. While it is hardly news that doing the right thing can sometimes require significant personal sacrifice, the objector may press that there might well be cases where the psychological difficulty of making such a sacrifice is so great that the agent in question simply cannot bring

---

<sup>8</sup> It is worth noting that Kim et al. also point out that ethically bounded agents can avoid violating their obligations by recusing themselves instead of acting on their biases (Kim et al. 2015, 353). I will return to the question as to how much of a difference there is between Kim et al.’s position and mine at the end of section 6.

themselves to do it. If this is granted, it can serve as the first premise in an argument analogous to RBC that yields the conclusion that the agent has no obligation to give up their career. Call this argument RBC\*.<sup>9</sup> It is an interesting question in its own right whether such an argument could succeed (it certainly seems even more counterintuitive than RBC, since it basically amounts to saying that a sufficient lack of motivation to fulfill one's obligations would negate their very existence). In the current context, however, it should suffice to point out that the psychological difficulty in question in such cases is different from the psychological difficulty presented by bounded ethicality in at least one important respect: part of the problem with bounded ethicality is that it affects our decision making without us ever becoming aware of this. This is not the case in situations where we cannot bring ourselves to sacrifice self-interest in order to do the right thing. Being aware of what we are doing and why is commonly thought to be a necessary condition for being morally responsible for those actions (Levy 2014; Holroyd et al. 2017). Thus, RBC and RBC\* do not stand or fall together.

## **6. Implications of the Wide-Scope View**

The suggestion that ethically bounded agents should simply recuse themselves from situations in which this is the only way to fulfill their wide-scope obligations would be problematic, if bounded ethicality affected everyone, always, and without remedy. For if this was true, nobody could ethically take on roles that clearly need to be filled in society. We would be forced to say, to stay with my main example, that every accountant is in violation of their ethical obligations. This would follow from Kim et al.'s proposed solution as well. Thus, they take pains to point out that, while the findings of bounded ethicality are robust and affect large majorities, there are, in every study,

---

<sup>9</sup> I would like to thank an anonymous referee for raising this objection.

some individuals who appear to be unaffected by these biases (Kim et al. 2015, 349; cf Greene 2018).

There are a number of different ways of thinking about what might explain this. The explanation that Kim et al. focus on is that people are, in principle, capable of fighting these nefarious influences (Kim et al. 2015, 349 and 355). If so, it is not true that we are psychologically incapable of acting in accordance with the prescriptions of normative business ethics. It may require some extra effort, but we knew that doing the right thing can sometimes be difficult long before anyone started to talk about heuristics, biases, and bounded ethicality. Thus, this is a puzzling stance for Kim et al. to take, for it makes their re-interpretation of OIC superfluous: if actual people can do what normative business ethics prescribes, we need not appeal to idealized rational agents. More importantly, however, it is doubtful that the difference between those affected and those unaffected by bounded ethicality in any particular circumstance is entirely a result of the latter trying harder to fulfil their obligations. It might explain some of the difference and, insofar as it does, the results of behavioural ethics give us reason to be more watchful whether we are about to unwittingly violate our obligations. But this cannot be the entire story.

In fact, researchers routinely warn against the strategy of ‘trying harder’ (Follenfant & Ric 2010; Soll et al. 2015). Many heuristics and biases tend to resist educational efforts, and even people who spend large parts of their lives studying these phenomena are subject to them. It is therefore largely futile to simply resolve to fight against one’s own bounded ethicality. A better strategy is to put institutional safeguards in place that insulate important decisions from our subconscious biases (cf Buckwalter 2018, 2967). Thus, in deciding whether or not to approve a loan application, it may be best to strictly follow the recommendations of a pre-determined

algorithm instead of trying to ‘get a feel’ for whether the applicant can be trusted.<sup>10</sup> An analogy might be helpful here. Imagine that, after encountering the Müller-Lyer illusion for the first time, you resolved to look very carefully, in the future, at the relative length of arrowed lines. This would be futile for, as we know, the illusion persists regardless of how hard you look and even after you are thoroughly satisfied intellectually that the lines are indeed the same length. The better strategy would be to never allow yourself to rely on your eyes in judging the relative length of arrowed lines. You should just institute a policy of actually measuring them (at least in situations in which getting it right matters). Finally, it is worth pointing out that there are bound to be individual differences in how susceptible we are to nefarious biases, and also in what kind of strategies work for each of us in fighting them (Stanovich and West 2008; Toplak et al. 2011).

The preceding is the merest sketch of the complexities that emerge once we drop the pretense, implicit in the formulation of the argument for RBC, that bounded ethicality affects everyone, always, and to the same degree. What I want to emphasize, however, is that the response I have recommended in section 5 is well equipped to do justice to these complexities. We all have a wide-scope obligation to ‘produce objective auditing reports, if taking on the role of auditor’. Just how we can fulfil this obligation depends on our individual psychological makeup, and circumstances. Some may be lucky enough to be simply able to do an unbiased audit; they ought to do so. Others might be able to do it with great mental effort, or by putting institutional safeguards in place; they ought to make sure to take the required precautions. And some may literally be

---

<sup>10</sup> This is not to advocate for placing blind trust in algorithms. As many have pointed out, algorithmic decision making comes with its own significant ethical pitfalls (see, e.g., O’Neil 2016). The point is to take the biases and limitations of human decision makers seriously, and install carefully designed institutional supplements. Such supplements need to be subject to rigorous and ongoing review.

incapable of producing an unbiased report. The ethical thing for them is to stay out of the accounting profession.

If all of this appears to presuppose an inordinate amount of self-knowledge, I would humbly suggest that we do have an obligation to acquire more self-knowledge than we might have thought sufficient before we learned about bounded ethicality. This is an important point to make. Recall that part of the reason why RBC is more plausible than RBC\* is that ethically bounded agents will often act contrary to their professed moral convictions without becoming aware that they are doing so. This very feature might seem to threaten the effectiveness of my solution to RBC. For, if people are faultlessly unaware of their transgressions, it is not clear that they have the occurrent ability to avoid them.<sup>11</sup> I think that this may very well have been the situation of ethically bounded agents before knowledge of bounded ethicality became widely available. No longer. As Washington and Kelly have argued, as soon as knowledge about bounded ethicality becomes available in the epistemic environment of an agent who occupies a social role to which such knowledge is relevant, it behooves that agent to acquire that knowledge (Washington and Kelly 2016). For example, depending on our evolving understanding of implicit bias, it might be plausible that anyone involved in hiring decisions ought to take an implicit associations test to determine whether they are at risk of unwittingly discriminating against certain groups.

That said, however, it should be clear that we cannot expect individuals to shoulder the entire burden of designing systems that insulate their decision making from their bounded ethicality. Many institutional safeguards are best erected and maintained by, well, institutions (cf Moore 2012). Since part of what makes bounded ethicality such an insidious phenomenon is that people behave unethically without ever becoming aware that their behaviour is even ethically

---

<sup>11</sup> I would like to thank the editor for raising this point.

significant in any way, it is important to create institutional knowledge of the ways in which things can go awry, so that individuals can be duly alerted and, where possible and necessary, trained. As Marcus put it 40 years ago: “we ought to conduct our lives and arrange our institutions so as to minimize predicaments of moral conflict” (Marcus 1980, 121).

As a simple example, there is clear evidence (a) that medical doctors’ prescription decisions are influenced by seemingly insignificant gifts from pharmaceutical companies, and (b) that MDs are blissfully unaware that this is the case (Lexchin 2017). Do MDs who succumb to these effects violate their obligation to prescribe with only their patients’ best interest in mind? Not if the RBC goes through, for then they would have no such obligation. According to Kim et al.’s suggestion, doctors do violate an obligation, but it is difficult to see what they could do about that other than to make an effort to fight tendencies they do not even know they have. By contrast, the suggestion here points to an institutional response. Medical associations ought to develop guidelines that take the research findings into account and ensure that the burden of fighting nefarious biases is not left with individual MDs many of whom might be literally incapable of doing so successfully.<sup>12</sup>

At this point, it is worth acknowledging that these recommendations are quite similar to the ones that Kim et al. make in the conclusion of their paper<sup>13</sup>:

Naturally, the knowledge that bounded ethicality exists may generate new obligations for moral actors—the obligation, for example, to undertake effective means to combat the relevant social and psychological influences of bounded ethicality. It is a significant

---

<sup>12</sup> Note that institutional solutions need to be devised and implemented by individuals, and can work only with sufficient buy-in from most individuals inside the institution. Thus, the suggestion is not to displace all responsibility from the individual to the institutional level, but rather that there is a complex interplay between individual and institutional responsibilities (cf Madva 2016).

<sup>13</sup> Thanks to an anonymous reviewer who insisted on this point.

contribution of behavioral business ethics that it helps us to think about the necessity of recognizing these new obligations. (Kim et al. 2015, 354)

Kim et al.'s overall argumentative strategy appears to be that appealing to the idealized version of OIC helps us to identify our obligations correctly, while our psychological limitations mean that we need to look at creative solutions of how we can fulfill them (Kim et al. 2015, 353). As I have been arguing, the latter conclusion can be reached without appealing to a version of OIC that refers to the abilities of an idealized rational agent. Indeed, much of the discussion in Kim et al. leading up to the conclusion rejects the premise that actual human agents are incapable of fulfilling the obligations identified by normative business ethics (see Kim et al. 2015, 348-353). It is thus somewhat unclear why Kim et al. believe the appeal to the idealized version of OIC is necessary.

More importantly, it is not clear that the idealized version of OIC is actually compatible with the (plausible) recommendations put forward in the cited passage. Consider that an ideal rational agent can neither improve their reasoning abilities (they are already perfect), nor take any steps to rein in predictable cognitive mistakes (there are none). Thus, if Kim et al.'s interpretation of OIC is correct, we can rule out obligations to improve one's reasoning abilities or erect safeguards against likely mistakes by an argument analogous to RBC:

Premise 1: An ideal rational agent cannot improve their reasoning abilities or erect safeguards against likely cognitive mistakes.

Premise 2: Necessarily, nobody has an obligation to do what an ideal rational agent cannot do.

Conclusion: Nobody has an obligation to improve their reasoning abilities or erect safeguards against likely cognitive mistakes.

This is an odd implication that Kim et al. would clearly want to avoid. Assuming that they are committed to premise 2 (their interpretation of OIC), they would have to reject premise 1. Thus,

we return to a point I made when discussing *MathBomb*: for Kim et al.'s solution to avoid counterintuitive implications, we would need to know much more about exactly what kind of capacities they take an idealized rational agent to have.

## **7. Towards Practical Advice**

Recall that my main worry about interpreting OIC as referring to the abilities of an idealized agent was that it risks turning normative business ethics into an endeavour that can merely condemn behaviour as unethical without offering advice as to how to avoid it. The remarks in the previous section point towards the way in which my proposed answer to RBC does better. In each concrete context, we need to ask how our limitations affect our ability to fulfill our obligations. It is likely that quite often the skeptic is right at the surface level: we will indeed find that we (or at least many of us) *cannot* comply with what we would intuitively think our obligations are. We should treat this as an invitation to take a closer look at what form our obligations really take, and what ways of meeting them there are that we *can* pursue. This is the right question to ask. And it is the question pursued by many of the researchers who made us aware of the challenge in the first place. To pick just a few, more or less random, examples, Kern and Chugh found that subjects are less likely to succumb to bounded ethicality when decisions are presented in a gain, rather than a loss, frame (Kern and Chugh 2009); Schweitzer et al. found that unintended unethical behaviour increases when very specific and demanding goals are being pursued (Schweitzer et al. 2004); Zhang et al. recommend that increasing subjects' vigilance at critical moments can make an important difference (Zhang et al. 2015); and Soll et al. present a whole battery of strategies for debiasing our decision making (Soll et al. 2015).

While here is not the place to go into detail about this research, the view I have recommended tells us why it is important: it helps us to figure out what we can do to fulfill our obligations in the face of the psychological limitations that we have recently become aware of. Insights about the presence of bounded ethicality and the mechanisms by which it operates can be turned into strategies for insulating our decision making from it.<sup>14</sup> The business application where this approach has been most progressively employed is probably hiring and employee review. Starting from the observation that women and members of minority groups are grossly underrepresented in many desirable occupations and leadership positions, researchers have developed a battery of interventions designed to eliminate bias in hiring and promotion decisions. Many of these interventions are quite simple, such as the decision to audition musicians for orchestra positions from behind a screen which led to a dramatic increase in the proportion of such positions going to women (Goldin and Rouse 2000). Somewhat more burdensome to implement, yet well worth doing, is the suggestion to eliminate unstructured interviews as well as panel interviews from the hiring process. Unstructured interviews give free reign to our unconscious biases and other forms of bounded ethicality while panel interviews amplify these effects by encouraging various forms of groupthink (Bohnet 2016, 139-141). It has also been shown that unconscious reliance on stereotypes can be significantly reduced by explicitly comparing candidates with each other, rather than trying to evaluate the merit of an individual candidate on its own (Bohnet 2016, 127).

---

<sup>14</sup> For a helpful framework for thinking about organizational design that helps to eliminate or insulate the foibles of individual members of an organization see de Bruin's discussion of strategies by which financial organizations can become more epistemically virtuous. De Bruin distinguishes three broad categories of such strategies: virtue-to-function matching, organizational support for virtue, and organizational remedies against vice (de Bruin 2015, 106-139).

It is important to note that erecting institutional safeguards against bounded ethicality is a delicate process that can easily go awry when insufficient attention is paid to whether interventions actually have the desired effects. In a recent review, van Rooij and Fine probe the efficacy of a trio of strategies to combat corporate crime: compliance management and ethics programs, whistleblower protection rules, and independent internal monitoring (van Rooij and Fine 2020). These three have all gained increasing prominence since the 1990s partly because the law is encouraging their use by, for example, looking more leniently at corporations that implement compliance and ethics programs (a provision of the Sarbanes-Oxley act). What van Rooij and Fine find, however, is that there is scant evidence that these strategies work. Their paper concludes with a call for more research and more detailed attention to the results of it when creating compliance strategies:

This means, first of all, using every available insight about how punishment, social and personal norms, capacity, opportunity, and unconscious influences can be employed in these particular situations. We need to change the training of compliance professionals, with less focus on studying legislation, court cases and legal procedures, and more on the criminology, psychology, sociology, and organizational science of how humans respond to rules. These insights must be translated to fit corporate organizational settings. (van Rooij and Fine 2019, 242)

Moreover, while some applications of the relevant research might be useful in a wide variety of settings, one must not simply assume that what works in one context will work in another as well. Rather it is important to pay attention to the details of the particular situation one is trying to address and create a mechanism that provides feedback as to whether any interventions are working as planned (Bohnet 2016, 13-14).

## **Conclusion**

Our increasing knowledge about the ways that our decision making is not fully controlled by our conscious beliefs and attitudes raises the prospect of a new skeptical challenge to normative business ethics. If we cannot follow the prescriptions of normative business ethics because of our psychological limitations, then it is not the case that we ought to follow these prescriptions. Accepting this claim amounts to throwing up one's hands, declaring that, in the face of human limitations, there is no way of bringing what is right, just, and fair into the business world. This would be an unacceptably defeatist attitude. It is, however, not a helpful response to simply insist that we ought to do the things that, as we have learned, we often cannot. The more helpful response, outlined in this paper, is to acknowledge the real difficulties that our psychological makeup poses to our ability to abide by our obligations, explicitly conceptualize those obligations as the complex, wide-scope principles they are, and proceed to explore creative ways in which we can fulfill them.

## **Acknowledgements**

Previous versions of this paper were presented during the 4th Workshop on Business Ethics at the European Institute for Advanced Studies in Management in Brussels, the Canadian Philosophical Association in Montreal, and the Society for Business Ethics in Chicago. I would like to thank the audiences at these occasions for valuable feedback, and would especially like to thank Santiago Mejia for detailed comments. A very early version of the piece profited greatly from feedback by Bruce Barry and Abraham Singer.

## References:

- Bertrand, M., Chugh, D., & Mullainathan, S. (2005). Implicit discrimination. *American Economic Review*, 94-98.
- Bohnet, I. (2016). *What works*. Harvard University Press.
- Brownstein, M., Madva, A., & Gawronski, B. (2020). Understanding implicit bias: Putting the criticism into perspective. *Pacific Philosophical Quarterly*, 101, 276-307.
- Buckwalter, W. (2019). Implicit attitudes and the ability argument. *Philosophical Studies*, 176(11), 2961-2990.
- Caruso, E. M., Epley, N., & Bazerman M. H. (2006). The costs and benefits of undoing egocentric responsibility assessments in groups. *Journal of personality and social psychology*, 91(5), 857-871.
- Chugh, D., Bazerman, M., Banaji, M. (2005). Bounded ethicality as a psychological barrier to recognizing conflicts of interest. *Conflicts of interest: Challenges and solutions in business, law, medicine, and public policy*, 74-95.
- De Bruin, B. (2015). *Ethics and the financial crisis: why incompetence is worse than greed*. Cambridge University Press.
- De Bruin, B. (2019). Epistemic injustice in finance. *Topoi*, 1-9.
- Feldman, Y. (2018). *The law of good people: Challenging states' ability to regulate human behavior*. Cambridge University Press.
- Follenfant, A., & Ric, F. (2010). Behavioral rebound following stereotype suppression. *European journal of social psychology*, 40(5), 774-782.

- Forscher, P. S., Lai, C. K., Axt, J. R., Ebersole, C. R., Herman, M., Devine, P. G., & Nosek, B. A. (2019). A meta-analysis of procedures to change implicit measures. *Journal of Personality and Social Psychology*, *117*(3), 522.
- Fricke, M. (2007). *Epistemic injustice*. Oxford University Press.
- Gawronski, B., Morrison, M., Phillips, C. E., & Galdi, S. (2017). Temporal stability of implicit and explicit measures: A longitudinal analysis. *Personality and Social Psychology Bulletin*, *43*(3), 300-312.
- Goldin, C., & Rouse, C. (2000). Orchestrating impartiality: The impact of "blind" auditions on female musicians. *American Economic Review*, *90*(4), 715-741.
- Graham, P. A. (2011). 'Ought' and Ability. *Philosophical Review*, *120*(3), 337-382.
- Greene, C. (2018). Mind the Gap: Virtue Ethics and the Financial Crisis. *Midwest Studies In Philosophy*, *42*, 174-190.
- Holroyd, J. (2012). Responsibility for implicit bias. *Journal of Social Philosophy*, *43*(3), 274-306.
- Holroyd, J., Scaife, R., & Stafford, T. (2017). Responsibility for implicit bias. *Philosophy Compass*, *12*(3), e12410.
- Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: a failure to disagree. *American psychologist*, *64*(6), 515.
- Kern, M. C., & Chugh, D. (2009). Bounded ethicality: The perils of loss framing. *Psychological Science*, *20*(3), 378-384.
- Kim, T., Monge, R., Strudler, A. (2015). Bounded ethicality and the principle that "ought" implies "can". *Business Ethics Quarterly* *25*(3), 341-361.

- Koehn, D. (2020). A virtue ethics critique of ethical dimensions of behavioral economics. *Business and Society Review*, 125 241-260.
- Kolodny, N. (2005). Why be rational?. *Mind* 114(455): 509-563.
- Lexchin, J. (2017). *Doctors in Denial*. Toronto: Lorimer and Co.
- Levy, N. (2014). Consciousness, implicit attitudes and moral responsibility. *Noûs*, 48(1), 21-40.
- Levy, N. (2017). Implicit bias and moral responsibility: Probing the data. *Philosophy and Phenomenological Research*, 94(1).
- Madva, A. (2016). A plea for anti-anti-individualism: How oversimple psychology misleads social policy. *Ergo, an Open Access Journal of Philosophy*, 3.
- Marcus, R. B. (1980). Moral dilemmas and consistency. *The Journal of Philosophy*, 77(3), 121-136.
- Messick, D., Sentis, K. (1979). Fairness and preference. *Journal of Experimental Social Psychology*, 15(4), 418-434.
- Moore, D., Tanlu, L., Bazerman, M. (2010). Conflict of interest and the intrusion of bias. *Judgment and Decision Making* 5(1), 37.
- Moore, G. (2012). The virtue of governance, the governance of virtue. *Business Ethics Quarterly*, 293-318.
- O’Neil, C. (2016). *Weapons of Math Destruction*. New York: Broadway Books.
- Oswald, F. L., Mitchell, G., Blanton, H., Jaccard, J., & Tetlock, P. E. (2013). Predicting ethnic and racial discrimination: a meta-analysis of IAT criterion studies. *Journal of personality and social psychology*, 105(2), 171.
- Pronin, E., Gilovich, T., & Ross, L. (2004). Objectivity in the eye of the beholder: divergent perceptions of bias in self versus others. *Psychological Review*, 111(3), 781-799.

- Rachlinski, J. J., Johnson, S. L., Wistrich, A. J., & Guthrie, C. (2008). Does unconscious racial bias affect trial judges. *Notre Dame L. Rev.*, 84, 1195.
- Saul, J. (2013). Scepticism and implicit bias. *Disputatio*, 5(37), 243-263.
- Schweitzer, M. E., Ordóñez, L., & Douma, B. (2004). Goal setting as a motivator of unethical behavior. *Academy of Management Journal*, 47(3), 422-432.
- Simon, H. (1957). *Models of Man*. New York: John Wiley
- Soll, J. B., Milkman, K. L., & Payne, J. W. (2015). A user's guide to debiasing. *The Wiley Blackwell handbook of judgment and decision making*, 2, 924-951.
- Stanovich, K. E., & West, R. F. (2008). On the Relative Independence of Thinking Biases and Cognitive Ability. *Journal of Personality and Social Psychology*, 94(4), 672-695.
- Stern, R. (2004). Does 'ought' imply 'can'? And did Kant think it does?. *Utilitas*, 16(1), 42-61.
- Toplak, M. E., West, R. F., & Stanovich, K. E. (2011). The Cognitive Reflection Test as a predictor of performance on heuristics-and-biases tasks. *Memory & cognition*, 39(7), 1275-1289.
- van Rooij, B., & Fine, A. D. (2019). Preventing Corporate Crime from Within: Compliance Management, Whistleblowing, and Internal Monitoring. *The Handbook of White-Collar Crime*, 229-245. Wiley Blackwell.
- von Kriegstein, H. (2019). Oxymoron: taking business ethics denial seriously. *Journal of Business Ethics Education*, 16, 103-134.
- Vranas, P. (2007). I ought, therefore I can. *Philosophical studies*, 136(2), 167-216.
- Washington, N., & Kelly, D. (2016). Who's responsible for this? Moral responsibility,

externalism, and knowledge about implicit bias. In Brownstein, M & Saul, J. (eds.)

*Implicit bias and philosophy: moral responsibility, structural injustice, and ethics* (vol 2, 11-36). OUP

Zhang, T., Fletcher, P. O., Gino, F., & Bazerman, M. H. (2015). Reducing bounded ethicality: How to help individuals notice and avoid unethical behavior. *Organizational Dynamics*, 44(4), 310-317.