

Ignorance and Duty: The Objective/Subjective Distinction in Ethics

Amelia Hicks

Abstract: The concept of ‘subjective’ moral obligation promises to help us understand the moral duties of ignorant agents. More specifically, the distinction between ‘objective’ and ‘subjective’ moral obligation is motivated by *dueling intuitions about the deontic statuses of actions*, by *action-guidance as a desideratum for a moral theory*, and by *intuitions about culpability*. However, this chapter argues that those three considerations fail to adequately motivate the objective/subjective distinction in ethics. The chapter ends with a discussion of how we can more fruitfully theorize about the moral relevance of ignorance.

Introduction

Consider a case in which an agent’s ignorance--their false beliefs, or their failure to hold true beliefs--seems to *make a moral difference*:

Accidental Orchid: Billie mistakenly believes that Alex loves orchids. So, Billie gives an orchid to Alex. In fact, orchids remind Alex of a terrible past experience. Alex is distressed to receive one as a gift.

It’s plausible that Billie’s mistaken belief makes *some* kind of moral difference. But is Billie’s mistaken belief relevant to the *moral deontic status*¹ of Billie’s action? One might think not. For example, according to a simple version of act consequentialism, the wrongness of an action is determined by the total utility of the action’s consequences, not by the agent’s beliefs. An act

¹ Moral deontic statuses include: *permissible*, *impermissible*, *obligatory*, and *supererogatory*.

consequentialist might say that Billie’s mistaken belief is morally relevant insofar as it *exculpates* Billie for their wrongdoing, and yet insist that Billie’s *action* remains wrong.²

However, this consequentialism-inspired view is incompatible with deep-seated intuitions that an agent’s beliefs *do* sometimes affect the deontic status of the agent’s action. For example, sometimes it seems like an agent who acts badly *could not have acted better*, given the beliefs they held. If ‘ought’ implies ‘can,’ and if an agent’s mistaken beliefs constrain what an agent can do, then an agent’s beliefs sometimes affect the deontic statuses of the agent’s actions.³ Thus, before we adopt the consequentialism-inspired view, we need a more persuasive story about *why*, exactly, mistaken beliefs never affect the deontic statuses of actions.

Perhaps Billie’s mistaken belief *does* affect the deontic status of Billie’s action. According to this view, Billie’s mistaken belief *makes* Billie’s action permissible, even though the action would have been impermissible had Billie believed the truth. The most extreme version of this view says that an action is permissible when it’s permitted according to the agent’s beliefs; however, even if this view yields a plausible verdict about Billie’s action in *Accidental Orchid*, there are good reasons to reject such an extreme form of moral subjectivism. For example, it’s implausible that one can make an impermissible action permissible by carelessly forming a false belief. Matters get even worse when we consider agents who hold false *moral* beliefs. Consider:

Intentional Orchid: Billie mistakenly believes in “tough love”--Billie thinks that one morally ought to confront people with their worst fears in order to toughen them up.

² Smart (1956) and Mill (1863) each propose their own version of this general idea. For a range of views about ignorance and exculpation, see: Harman (2011); Rosen (2004); Sliwa (2019); Smith (2015); Zimmerman (1997).

³ Hicks (2022).

Billie is aware that orchids are distressing for Alex, but delivers an orchid to Alex anyway--after all, Billie wants Alex to toughen up.

If an action is permissible whenever it's permissible according to the agent's beliefs, then Billie's action in Intentional Orchid is permissible. But in this case, Billie's action doesn't seem permissible. More generally, we ought not endorse the principle that actions resulting from false moral beliefs are always permitted.

Thus, if we think that mistaken beliefs *can* affect the wrongness of actions, and we want to avoid an implausibly permissive version of moral subjectivism, then we must draw a line between mistaken beliefs that can affect the deontic status of an action and those that cannot. Perhaps we could say that only *certain* false beliefs--false beliefs with the "right type" of epistemic history--affect the deontic status of actions. However, it's plausible that horrifying false beliefs can have respectable epistemic histories,⁴ and so we cannot avoid the problem of excessive permissiveness by insisting that only beliefs with respectable epistemic histories affect the deontic status of actions.

When it comes to the deontic status of Billie's action in Accidental Orchid, we seem to be stuck. On the one hand, it's plausible that Billie's mistaken belief could, in principle, affect the deontic status of Billie's action. On the other hand, it's implausible that an action is permissible whenever it's permitted according to the agent's beliefs--such a view is too permissive.

⁴ Someone who lives in an environment in which they lack access to high quality evidence might do their best to form beliefs carefully, and yet still end up holding horrifying false beliefs.

Moreover, it's unclear how we can draw a well-motivated line between ignorance that affects the deontic status of actions and ignorance that doesn't.

One strategy for getting ourselves unstuck is to refrain from trying to determine the conditions under which ignorance affects the deontic status of actions. Instead, we can introduce a distinction between two types of moral obligation: *objective* moral obligation and *subjective* moral obligation. One's objective moral obligation is determined by all of the descriptive and moral facts, *not* by one's beliefs; one's subjective moral obligation is determined at least in part by one's beliefs.⁵ Perhaps our confusion about the moral relevance of Billie's beliefs is traceable to confusion about the very concept of moral obligation; we have been trying to provide a moral analysis of Billie's action using only a single concept of moral obligation, when really we need *two* such concepts. Perhaps Billie's mistaken belief in *Accidental Orchid* made their action *subjectively* permissible, even though Billie's action remained *objectively* impermissible.

This chapter critically examines the distinction between objective and subjective moral obligation. Section 1 explores common motivations for the distinction, and Section 2 argues that many of those motivations are suspect. Sections 1 and 2 work with an intuitive characterization of the objective/subjective distinction, according to which objective moral obligations are insensitive to agents' beliefs, and subjective moral obligations are sensitive to agents' beliefs. Section 3 more carefully examines different characterizations of the distinction, and challenges the distinction's coherence. Finally, Section 4 suggests an alternative approach to understanding the moral relevance of ignorance.

⁵ There are many different characterizations of the subjective/objective moral obligation distinction. See: Feldman (2012); Graham (2010); Jackson (1991); Olsen (2017); Smith (2010); Zimmerman (2006).

This chapter does *not* argue that there is only a single concept of “obligation”; for example, it’s reasonable to think that there exist both *rational* and *moral* obligations. Moreover, the objective/subjective distinction might make sense in some *non-moral* normative domains (e.g., within the domain of epistemic normativity). But as we will see, when it comes to *moral* obligations, we have good reason to resist drawing the objective/subjective distinction.

1. Roles of the Objective/Subjective Distinction in Ethics

Ethicists often introduce the objective/subjective obligation distinction to remedy conceptual confusion. According to those who endorse the distinction, there are (at least) two distinct types of “moral obligation,” and thus moral analyses that employ only a single concept of moral obligation are inadequate. There are three main motivations for the distinction: *dueling intuitions about the deontic statuses of actions*, *action-guidance as a desideratum for a moral theory*, and *intuitions about blameless wrongdoing*.⁶⁷

1.1 Dueling Intuitions

Agents often act out of ignorance; and in some of these cases, our moral intuitions pull us in different directions. For example, Billie’s action in *Accidental Orchid* seems wrong in one sense (because it causes Alex significant distress), and seems permissible in another sense (because Billie sincerely believes they are giving a nice gift). More famously, a physician who administers a harmful drug on the basis of excellent evidence seems to act both rightly and wrongly at the

⁶ Smith (2010) also identifies these three considerations as motivating the distinction.

⁷ Some philosophers draw the distinction, only to then embrace one of the conceptions of moral obligation as the “real one.” Because these philosophers do not think there are *multiple* types of moral obligation, this chapter doesn’t discuss their views at length. For example, see Graham (2010).

same time;⁸ intuitively, they act rightly in the sense that they do the best they can given their beliefs and evidence, and yet they act wrongly in the sense that they harm their patient. Moreover, there are cases in which, intuitively, one subjectively ought to perform an action one *knows* to be objectively wrong. For example, you know that 100 miners are trapped in a mineshaft--but you have no idea whether they're in shaft A or shaft B. As a matter of fact, they're in shaft A. If you do nothing, then both shafts will partially flood and 10 miners will die; if you block shaft A, then no miners will die; and if you block shaft B, then all of the miners will die. In this case, intuitively, you objectively ought to block shaft A, because doing so would save all of the miners. But blocking shaft A would be extremely *risky* (given your evidence, you would run a 50% chance of killing all of the miners), and so *subjectively* you ought to do nothing--it's better to save 90 miners with certainty than to run a 50% risk of killing all of the miners.⁹

These are all examples of *dueling intuitions*--each pair of intuitions pulls us toward two distinct moral analyses of actions. It's natural, then, to suspect that we are operating with two different conceptions of moral obligation: one according to which an agent's beliefs are morally relevant (subjective obligation) and one according to which an agent's beliefs are not morally relevant (objective obligation). With the subjective/objective distinction in place, we can say that Billie's action in *Accidental Orchid* is both permissible and impermissible, albeit in two different senses.

⁸ Jackson (1991).

⁹ Regan (1980): p. 265, footnote 1.

1.2 Action-Guidance

Many philosophers treat *action-guidance* as a desideratum for a moral theory; this desideratum, in turn, motivates the distinction between objective and subjective moral obligation. One does not always have *epistemic access* to one's objective moral obligations, and thus one cannot reliably "consult" one's objective moral obligations in order to decide how to act. Subjective obligations, by contrast, are determined (at least in part) by one's beliefs and credences, and thus one more reliably retains epistemic access to one's subjective obligations. Because subjective obligations are more epistemically accessible, they are helpful, action-guiding supplements to one's objective moral obligations.¹⁰

There are at least three ways in which an objective moral norm can fail to guide action. To illustrate each of these possible failures, let's work with a plausible objective moral obligation: *one ought not kill sentient beings without a weighty reason*. The first way in which this norm can fail to guide action is if the agent fails to hold true descriptive beliefs, and as a result cannot accurately *apply* the norm. For example, one can know that one ought not kill sentient beings without a weighty reason, and yet one might also hold the false descriptive belief that octopuses are not sentient. In this type of case, one could end up falsely believing that it's permissible to kill octopuses without a weighty reason, and thus one will struggle to apply this norm when it comes to octopuses. Second, *moral ignorance* prevents objective moral norms from reliably guiding action. Sure, it might be objectively true that one ought not kill sentient beings without a weighty reason; but if someone is raised in a society in which it's routine to kill sentient beings for no good reason, then that person will probably not learn the objective moral truth. Such a

¹⁰ Smith (2010), Smith (2018).

person might know that octopuses are sentient, and yet kill octopuses anyway because they're unaware of the objective norm that forbids the frivolous killing of sentient beings.

There is a third way in which an objective norm can fail to guide action. Imagine that an agent knows the objective moral norm that *frivolously killing sentient beings is forbidden*; and imagine that they also hold correct descriptive beliefs about which beings are sentient. Still, the agent might fail to know *that they know* the correct moral norm, or they might fail to know *that they know* which beings are sentient. (Imagine someone who knows that *frivolously killing sentient beings is forbidden*, but who doubts their own knowledge because of social pressures from their local culture.) These failures of second-order knowledge can (if reflected upon by the agent) cause the agent to experience a peculiar type of uncertainty; and that uncertainty, in turn, threatens the action-guiding power of the objective moral norm. If one is uncertain about the correctness of one's beliefs, one might reasonably hesitate to act on those beliefs, and one might instead look to principles that guide action under conditions of uncertainty.¹¹

Subjective moral norms--which are partly determined by an agent's beliefs, or at least by the information an agent has access to--are more reliably action-guiding. For example, according to Holly Smith, what one subjectively ought to do is determined by the objective moral truths alongside one's descriptive beliefs; returning to our earlier example, Smith would say that one is subjectively permitted to kill an octopus if one falsely believes that octopuses are not sentient.¹² Thus, an agent who knows that it's wrong to frivolously kill sentient beings but who fails to know which beings are sentient can satisfy a subjective moral norm even as they violate an

¹¹ For an overview of some candidate principles, see MacAskill, Bykvist, and Ord (2020).

¹² Smith (2010), Smith (2018).

objective moral norm. Subjective moral norms can avoid the first type of failure of action-guidance, because the descriptively mistaken agent can use the subjective moral norm to chart a course for action if they want to.¹³

Consequentialists sometimes adopt the idea that subjective moral norms are action-guiding supplements to objective moral norms.¹⁴ If one endorses a version of consequentialism according to which one ought to maximize actual good consequences, then one's moral theory is not reliably action-guiding; after all, one rarely knows which of one's available actions will maximize good consequences. But the consequentialist can consistently maintain that, for practical purposes, one ought to act in a way that maximizes good consequences *in expectation*, taking into account the information the agent knows (or has some reasonable credence in, or has epistemic access to). According to this type of consequentialism, objective moral norms are extremely difficult to satisfy (because of human epistemic limitations), but subjective moral norms help agents do their best given the information they have. Moreover, non-consequentialists can endorse the subjective/objective distinction for similar reasons. For example, it's coherent for a deontologist to distinguish between (a) what one ought to do given the moral law and all the morally relevant facts, and (b) what one ought to do given the moral law and one's descriptive *beliefs*.

Subjective moral norms are more action-guiding than objective moral norms, at least when it comes to agents who fail to hold true descriptive beliefs. However, as will be discussed in Section 2.2, subjective moral norms do not address the other two failures of action-guidance

¹³ Subjective moral norms satisfy Smith's definition of an "internal decision guide." See Smith (2010), p. 74. Thus, Smith would say that the second two failures of action-guidance are not *genuine* failures of action-guidance.

¹⁴ Feldman (2012), Railton (1984).

(failures that result from moral ignorance, and failures that result from uncertainty). The fact that even subjective moral norms fail to reliably guide action forces us to ask: should we lower the bar for when a moral theory counts as adequately action-guiding? If so, is a supplemental theory of subjective moral obligation still necessary for a moral theory to count as adequately action-guiding?

1.3 Blame

Finally, intuitions about blameless wrongdoing sometimes motivate the subjective/objective moral obligation distinction. Return to Billie in *Accidental Orchid*. Many of us intuit that Billie acts wrongly, and yet is blameless for their wrongdoing. But we might wonder *why* Billie's mistaken belief exculpates Billie for performing an impermissible action--is it just a brute fact about how blameworthiness works, or can we provide a deeper explanation for Billie's exculpation? One might suspect that the explanation for Billie's exculpation is Billie's satisfaction of a *subjective* moral norm.

According to this type of view, whether an agent satisfies an objective moral obligation determines the deontic status of the agent's action; and whether the agent satisfies a subjective moral obligation determines the agent's *culpability* for the action. Because it's possible for an agent to satisfy a subjective obligation while violating an objective obligation (like the physician who administers a harmful drug on the basis of excellent evidence), instances of blameless wrongdoing are possible; and because it's possible for an agent to satisfy an objective obligation while violating a subjective obligation (like a physician who tries to harm their patient, but who administers a helpful drug by accident), instances of culpable rightdoing are possible. On this

view, the objective/subjective distinction explains why (and how) the deontic status of an action can vary independently of the agent's culpability for that action.

2. Some Challenges

We've seen that the objective/subjective moral obligation distinction plays several roles in moral theorizing: it accounts for dueling intuitions about deontic status, it explains how a moral theory can provide action-guidance, and it explains intuitions about blameless wrongdoing. But does the distinction play these roles *well*?

2.1 Problems with Dueling Intuitions

There are three problems with appealing to dueling intuitions about deontic status to motivate the distinction between subjective and objective moral obligation. First, it's unclear whether those intuitions are reliable. Second, even if we grant that those intuitions are worth attending to, they might track a *different* distinction. And third, using dueling intuitions to motivate the objective/subjective distinction raises difficult questions about the relationship between the two concepts of moral obligation.

Intuitively, Billie's action in Accidental Orchid seems *permissible* in one sense and *impermissible* in another. But just how much are these intuitions worth? Our first-order moral intuitions are notoriously unreliable; the extent to which first-order moral intuitions diverge demonstrates that many of us have inaccurate moral intuitions. Moreover, the intuitions elicited by cases such as Accidental Orchid are particularly suspect, because they are dueling first-order moral intuitions (*Billie's action is permissible* and *Billie's action is impermissible*) that are being

used to support a *metaethical* thesis (namely, *that there are two distinct types of moral obligation*). Someone who embraces the subjective/objective distinction on the basis of dueling intuitions about deontic status must provide some assurance that these intuitions are trustworthy, and are not mere indicators of conceptual confusion.

But let's assume that these intuitions are tracking *something* that's worthy of in-depth moral analysis. Still, the distinction these intuitions are tracking need not be the objective/subjective moral obligation distinction. Most moral theories include at least the following two components: a *theory of deontic status* (an account of the conditions under which actions count as permissible, impermissible, obligatory, supererogatory, etc.), and an *axiological theory* (an account of what's morally valuable). It's conceivable we intuit that Billie's action in Accidental Orchid is permissible, because Billie's action *is* permissible; at the same time, it's conceivable we intuit that Billie's action in Accidental Orchid is impermissible, because we recognize that it would be *better* for Billie to perform a different action. In other words, our "permissible" intuition might track the deontic status of Billie's action, while our "impermissible" intuition might track an *axiological* fact about Billie's action. This is coherent, as long as we allow that an agent sometimes ought to perform a sub-optimal action. And it's plausible that agents sometimes ought to perform sub-optimal actions; after all, optimal actions are not always possible.¹⁵ Thus, our dueling intuitions about cases like Accidental Orchid may very well be tracking a distinction *other* than the objective/subjective distinction.

Let's now assume that our dueling intuitions are tracking the objective/subjective distinction. If our dueling intuitions motivate the distinction, then we're left with difficult questions about the

¹⁵ Hicks (2022).

relationship between subjective and objective moral obligation. Is one of these conceptions “primary”? When the two types of obligations diverge, which one should the agent *really* comply with? Do we need a third type of moral obligation that adjudicates conflicts between the subjective ‘ought’ and the objective ‘ought’? Notice that these difficulties don’t immediately arise when the distinction is motivated by intuitions about blameworthiness or action-guidance. If intuitions about blameworthiness motivate the distinction, then we can simply say that whether one satisfies an objective obligation determines the deontic status of one’s action, and at the same time we can say that whether one satisfies a subjective obligation determines one’s culpability--there’s no “competition” between two distinct oughts. When the objective/subjective distinction is motivated only by action-guidance considerations, we can simply say that one *really* ought to satisfy one’s objective obligations, but one may use one’s subjective obligations as a practical tool for approximating the satisfaction of one’s objective obligations--again, there’s not necessarily any competition between two distinct, equally weighty oughts. But if the distinction is motivated by dueling intuitions about deontic status, then we’re left wondering how objective and subjective moral obligation are related, and we’re left wondering about the structure of a moral theory that includes both.

2.2 Problems with Action-Guidance

Recall that objective moral norms suffer from three types of failures to guide action: they do not reliably guide the actions of agents who hold false descriptive beliefs, they do not reliably guide the actions of agents who hold false moral beliefs, and they do not reliably guide the actions of agents who second-guess their beliefs. Subjective moral norms are more successful at guiding

action; according to some popular accounts of subjective moral obligation,¹⁶ an agent can comply with a subjective moral obligation even if they hold false descriptive beliefs.

However, subjective moral norms cannot address the second two types of failures of action-guidance. Unless one adopts an extreme characterization of subjective moral obligation (according to which one is subjectively obligated to do whatever one believes one is obligated to do), subjective moral obligations cannot reliably guide the actions of morally ignorant agents. According to Smith, who endorses the objective/subjective distinction on the basis of considerations related to action-guidance, only descriptive ignorance (not moral ignorance) affects one's subjective moral obligations, because only descriptive ignorance impedes one's ability to use moral principles.¹⁷ But this means that a morally ignorant agent can reliably comply with neither their objective moral obligations nor their subjective moral obligations. Similarly, an agent who doubts their own beliefs--either their moral or descriptive beliefs--cannot reliably satisfy their subjective moral obligations.¹⁸

These observations demonstrate that if we reject an implausibly extreme characterization of subjective moral obligations (a characterization according to which one subjectively ought to do whatever one believes one ought to do), then subjective moral obligations end up being epistemically inaccessible in a number of different situations, much like objective moral obligations. When a moral norm is not reliably epistemically accessible, it cannot reliably be

¹⁶ Harman (2015); Smith (2010); Smith (2018).

¹⁷ Smith (2018) rejects the idea that moral ignorance undermines one's ability to use moral principles, because she treats moral ignorance as a lack of moral *motivation*. She writes, "[L]ack of motivation to follow a principle does not affect a person's ability to use the principle in the normal sense, any more than lack of motivation to do so affects a person's ability to use the can-opener in his kitchen drawer" (p. 17).

¹⁸ Smith (2010, 2018) would say that a morally ignorant agent (as well as an agent who second-guesses their beliefs) can still satisfy their subjective moral obligations, because the agent can use those subjective moral norms to guide their actions *if they want to*.

used as a guide for action. Thus, plausibly described subjective moral obligations are not reliable guides for action. Given that subjective moral obligations sometimes fail to guide action, we should explore whether action-guidance considerations adequately motivate the objective/subjective distinction.

At this point, one might point out that subjective moral obligations are clearly *more* action-guiding than objective moral obligations; a descriptively ignorant agent can comply with a subjective moral obligation even when they cannot (reliably) comply with an objective moral obligation. Perhaps it is only this additional type of action-guidance--action-guidance for the descriptively ignorant--that's needed to motivate the introduction of a new type of moral 'ought.'¹⁹ However, this suggestion (1) assumes a clear distinction between descriptive and moral beliefs, and (2) assumes that the role of a moral theory is to guide the actions of descriptively ignorant agents, but not morally ignorant agents.

We have reason to doubt that there is a hard-and-fast distinction between moral and descriptive beliefs. For example, some false descriptive beliefs are held as a result of moral failings. When someone believes *that women are cognitively inferior to men*, that false descriptive belief is likely the result of prior moral failings (including false moral beliefs)--false moral beliefs can affect the way in which an agent gathers evidence, which can in turn lead to the formation of false descriptive beliefs. In these sorts of cases, false descriptive beliefs take on a distinctively moral valence, even if the belief content is descriptive. Moreover, beliefs that employ thick moral concepts--such as the belief *that Annie is generous*--are in some respects moral beliefs but

¹⁹ Smith writes that this type of action-guidance is sufficient to "secure an agent's autonomy" (Smith thinks that a moral theory should be action-guiding so as to secure the autonomy of agents). Smith (2018), pp. 194-5.

in other respects descriptive beliefs.²⁰ This blurry boundary between moral and descriptive belief content creates problems for accounts of subjective moral obligation according to which one's subjective moral obligations are affected only by one's descriptive beliefs and not by one's moral beliefs.

Even if we set aside our worries about the boundary between moral and descriptive beliefs, it's still odd to demand that a moral theory reliably guide the actions of descriptively ignorant agents but *not* the actions of morally ignorant agents. It's natural to think that, among other things, a moral theory ought to be practical--it ought to provide us with tools for making moral decisions. But, of course, most of us are both descriptively ignorant (in some respects) *and* morally ignorant (in some respects). If we conceive of subjective moral obligations as obligations one has in light of the moral truth in combination with one's descriptive beliefs, and if we acknowledge just how epistemically and motivationally limited human beings are (and in particular, if we acknowledge that most of us never succeed in believing "the entire moral truth"), then it turns out that subjective moral obligations are not much more action-guiding than objective moral obligations. Only those who happen to already believe the moral truth can reliably comply with subjective moral obligations; the rest of us are no more guided by our subjective moral obligations than by our objective obligations.

So far, this subsection has suggested that most people experience at least some moral ignorance, and thus most people cannot be reliably guided by subjective moral norms. If this is correct, then considerations related to action-guidance do not adequately motivate the objective/subjective moral obligation distinction. However, many philosophers think that moral truths are

²⁰ Väyrynen (2021).

epistemically accessible in a way that descriptive truths are not. Some philosophers think that moral truths are a priori, and can be discovered through reason alone.²¹ Call this view *moral rationalism*. Moral rationalism addresses some of the worries just raised. In particular, if moral rationalism is true, then any rational agent can (in principle) come to believe the fundamental moral truths, just by thinking hard; and if any rational agent can come to believe the moral truth, then any rational agent can, in principle, comply with their subjective moral obligations (even if they can't reliably come to believe descriptive, empirical truths).

This is not the place to adjudicate moral rationalism, although there are reasons to doubt it.²² Still, we can summarize the foregoing by highlighting some background assumptions one must make if one wants to use action-guidance considerations to motivate drawing a distinction between subjective and objective moral obligation. First, if one *rejects* moral rationalism, then one must admit that it's often impossible for an agent to discover the moral truth. In that case, the only additional form of action-guidance a subjective moral norm can provide is action-guidance for the merely descriptively ignorant. So, if one rejects rationalism and wants to introduce the subjective/objective obligation distinction so as to increase the action-guiding power of a moral theory, then one must assume that the only type of action-guidance a moral theory should provide is guidance for the *descriptively* ignorant. This is an extremely contentious assumption, especially given reasonable pessimism about humans' ability to uncover all moral truths. Second, one can accept moral rationalism, and thereby maintain that subjective moral obligations (which are affected only by an agent's descriptive beliefs) are epistemically accessible, at least in

²¹ For recent work on moral rationalism, see Jones and Schroeter (2018).

²² Kauppinen (2018).

principle, to all rational agents. But this strategy requires the assumption that moral rationalism is true--also a very contentious assumption.

One might, at this point, be tempted to adopt a more extreme characterization of subjective moral obligation, according to which an agent's subjective moral obligations are determined by both their descriptive *and* their moral beliefs. By working with this more extreme characterization of subjective moral obligation, subjective moral obligations become reliably epistemically accessible (at least for agents who are capable of introspecting their own mental states).

Additionally, this view allows us to maintain that subjective moral norms are reliably accessible to agents *without* assuming moral rationalism. But recall just how extreme this type of view is.

According to this account of subjective moral obligation, a person who holds horrifying false moral beliefs turns out to be subjectively morally obligated to act on those beliefs.²³ Thus, on this view, subjective moral norms simply aren't very *interesting*, from a moral perspective; the only reason they're reliably discoverable by any moral agent is that they have no independent content whatsoever.

Thus, if we reject moral rationalism, we face a trade-off between a moral norm's *substance* and its *epistemic accessibility*.²⁴ A subjective moral norm that's reliably epistemically accessible--even by highly flawed agents--must lack substance (its only substance is supplied by the flawed beliefs of the agent); and its lack of substance prevents it from being a useful guide for action. In other words, an epistemically accessible subjective moral norm is no longer a helpful tool for approximating the satisfaction of objective moral norms. But a more substantive

²³ As will be discussed in section 2.3, we cannot avoid this result even if we place constraints on the epistemic histories of the beliefs that determine one's subjective moral obligations.

²⁴ Weatherson (2019) makes a similar point.

subjective moral norm will not be reliably epistemically accessible; so, more substantive subjective moral norms are not more practically useful than objective moral norms, at least not by much.

2.3 Problems with Blameworthiness

Recall that intuitions about blameworthiness sometimes motivate the distinction between subjective and objective moral obligation. Perhaps cases of blameless wrongdoing are cases in which the agent violates their objective obligation and satisfies their subjective obligation; and perhaps cases of blameworthy rightdoing are cases in which the agent satisfies their objective obligation but violates their subjective obligation. However, there are three disadvantages to using these intuitions to motivate the distinction. First, an agent's (non-)compliance with their subjective obligation doesn't perfectly track our intuitions about blameworthiness. Second, if we place constraints on the epistemic histories of the beliefs that can generate subjective moral obligations, then probably very few people have subjective moral obligations and very few people are ever exculpated for their objective wrongdoing. And third, the concept of blameworthiness is itself contested.

To see why an agent's compliance/non-compliance with their subjective obligations doesn't always track our intuitions about the agent's blameworthiness, assume a common conception of subjective moral obligation: one is subjectively obligated to perform the action that the objectively correct moral theory would identify as obligatory if the agent's descriptive beliefs were correct. On this conception of subjective moral obligation, there will be cases in which the agent violates their subjective moral obligation and yet is intuitively blameless. To see this,

imagine any scenario in which it seems like an agent blamelessly forms a false moral belief; we can imagine a version of Billie who has been systematically indoctrinated into believing in moral tough love, who has sincerely examined their commitment to moral tough love but (because of their indoctrination) cannot see its flaws. So, it seems like certain false moral beliefs can be formed non-culpably. And yet, if false moral beliefs cannot affect an agent's subjective moral obligations, and if the concept of subjective obligation is being used to explain the conditions under which an agent is blameworthy, then it turns out that agents who act on non-culpably formed false moral beliefs still violate their subjective moral obligation, and are rendered blameworthy for their objective wrongdoing.

Here is a different type of example, in which the agent is intuitively blameworthy for their objective wrongdoing, and yet using the concept of subjective obligation to determine blameworthiness instead yields the verdict that the agent is blameless. Consider a case in which an agent culpably lacks morally salient information. For example, think of a version of Billie in *Accidental Orchid* who is mistaken about Alex's flower preferences, because Billie wasn't listening when Alex revealed their extreme distaste for orchids. In this case, Billie's failure to figure out that Alex is distressed by orchids is the result of thoughtlessness. But if an agent's subjective moral obligation is determined by the moral truth in combination with the agent's descriptive beliefs, and if the concept of subjective moral obligation is being used to determine the agent's culpability for their action, then it turns out that this version of Billie acts blamelessly when they give Alex the distressing orchids. But presumably Billie is blameworthy, because their gift was the result of thoughtlessness.

There is a possible solution to the problems just described. We could say that only beliefs with the *right sort of epistemic histories* can affect an agent's subjective moral obligations. Given that Billie's belief that Alex likes orchids has a *thoughtless* epistemic history, it's not the type of belief that can generate a subjective moral obligation; so, when Billie acts on that belief, they are not exculpated for their objective wrongdoing. We can even extend this idea to moral beliefs; perhaps if a moral belief (even a false one) has the right type of epistemic history, then it can affect the agent's subjective moral obligations. Perhaps the version of Billie who believes in moral tough love as a result of indoctrination acts blamelessly, because of the epistemic history of that belief (assuming Billie wasn't culpable for their indoctrination).

Although it's reasonable to think that we ought to place historical epistemic constraints on the beliefs that determine subjective moral obligations, this solution to our earlier problems creates problems of its own. Let's say that according to this revised conception of subjective moral obligation, only beliefs with flawless epistemic histories determine an agent's subjective moral obligations. The first problem with this revision is that it seems like there are truly horrifying beliefs that could, in principle, have flawless epistemic histories. It seems possible for someone in an epistemically impoverished environment (an environment in which they lack access to high quality evidence) to reasonably come to hold horrifying beliefs. Arguably such beliefs have flawless epistemic histories, in the sense that the agent could not have reasonably believed otherwise under the circumstances; and yet, we hesitate to say that someone who acts on those horrifying beliefs acts blamelessly. Perhaps we should revise our conception of subjective obligation again, so that it says that an agent's subjective moral obligations are determined by the moral truth in combination with the agent's *descriptive* beliefs that have epistemically flawless

histories. This revision rules out the possibility of an agent's subjective obligations being affected by their horrifying moral beliefs. Still, we might worry about how many agents--or even whether *any* agents--have descriptive beliefs with truly 'flawless' epistemic histories. If we hold a pessimistic view about humans' cognitive capacities, then it's plausible that very few of our beliefs (if any) have been formed flawlessly. We can, of course, lower the standard for the epistemic histories of beliefs that generate subjective moral obligations; but then we must determine where, exactly, that standard lies. Moreover, if we lower the standard for the epistemic histories of beliefs, then we will almost certainly encounter more cases in which an agent is "exculpated" for their objective wrongdoing even though they are intuitively blameworthy.

Why is it so difficult to draw a conceptual link between subjective moral obligation and culpability? No doubt the difficulty in part stems from imprecision in how we've formulated the concept of subjective moral obligation--a problem we will turn to in the next section. But the difficulty also stems from unclarity about the concept of blameworthiness. There are many different conceptions of blameworthiness, and it's unclear which (if any) of those conceptions is linked to subjective moral obligation. For example, imagine an agent who flawlessly forms a horrifying false moral belief; given their circumstances and evidence, they were rationally compelled to adopt it. Such an agent might be blameless *in the accountability sense* for the actions that proceed from that belief, because they could not have reasonably believed otherwise. And yet, that agent might be blameworthy *in the attributability sense* for actions that proceed from that belief, because holding such a horrifying belief is constitutive of being a bad moral agent, and because the actions that proceed from that belief are attributable to the character of the agent.²⁵ We can, of course, attempt to define a particular type of blameworthiness that is

²⁵ Shoemaker (2011).

conceptually linked to subjective moral obligation. But in doing so, we must *ignore* a number of strong intuitions that stem from other conceptions of blameworthiness.

Thus, our intuitions about blameworthiness do not obviously support drawing a distinction between objective and subjective moral obligation. Our intuitions about an agent's blameworthiness do not cleanly track the agent's satisfaction of subjective moral obligations; introducing epistemic constraints on the beliefs that can generate subjective moral obligations solves some of those problems but creates others; and our intuitions about blameworthiness are all over the place, in part because there are so many different ways of conceptualizing blameworthiness.

3. The Incoherence of the Subjective/Objective Moral Obligation Distinction

So far, we have worked with a rough characterization of the objective/subjective moral obligation distinction. According to this rough characterization, one's objective moral obligations are determined by "all the moral and descriptive facts" and not by one's beliefs, whereas one's subjective moral obligation is determined, at least in part, by (some of) one's beliefs.

3.1 What's Objective Obligation?

According to the rough characterization of objective moral obligation, one's objective moral obligations are *belief-insensitive*; beliefs cannot affect the agent's objective moral obligations. But what does it mean to say that objective moral obligations are belief-insensitive? It might literally mean that beliefs never affect objective moral obligations, but such a view is

implausible. For example, imagine a scenario in which I see that my neighbor's house is on fire, and I know that my neighbors aren't home. Surely facts about my neighbors' belief-states--e.g., the fact that they don't believe that their house is on fire--*make it the case* that I ought to alert them about the fire (after calling the fire department, of course). Clearly, facts about beliefs can play a role in determining what an agent objectively ought to do.

Perhaps objective moral obligations are belief-insensitive in the sense that the *agent's* beliefs play no role in determining the agent's objective moral obligations. But even this view is implausible. Many philosophers (including some consequentialists²⁶) think that act-types are morally relevant; in other words, many philosophers think that whether an action falls under some type (such as *lying* or *killing*) affects the deontic status of that action. But an agent's beliefs play a role in determining the type under which their action falls. For example, imagine two people who utter the very same false assertion ("Fun fact: pineapples are actually vegetables!"). One person sincerely believes their assertion, while the other person disbelieves their assertion. The person who sincerely believes their assertion did not lie; but the person who disbelieves their assertion lied. The difference between the lie and the sincere utterance of a falsehood is the agent's *belief* in their own assertion. Thus, if we grant that an action's type can affect the deontic status of that action, then it turns out that an agent's beliefs can affect the deontic status of the action.

At this point, one might dig in one's heels, and insist that act-types are never morally relevant, thereby rescuing the idea that objective obligations are belief-insensitive (in the sense that an agent's beliefs never affect the agent's objective moral obligations). But an agent's beliefs can

²⁶ Hooker (2000).

affect the deontic status of that agent's actions in other ways. For example, it's plausible that an agent's beliefs affect which actions are *possible* for the agent to perform (in the morally relevant sense of "possible"). Assuming 'ought' implies 'can,' an agent's beliefs can affect which actions that agent ought to perform. Imagine a scenario in which Billie's nextdoor neighbor needs immediate medical attention. Although Billie is at home, Billie is completely unaware of their neighbor's needs (all is perfectly quiet). Sure, it would be physically possible for Billie to walk next door and assist the neighbor; but given Billie's beliefs, walking next door to provide medical assistance would *make no sense* to Billie. Walking next door would involve responding to moral reasons of which Billie is completely unaware. Because walking next door wouldn't, for Billie, be a way of responding to moral reasons, doing so wouldn't be an exercise of Billie's agency. We might say that walking next door to attend to the neighbor isn't *agentially possible* for Billie, because (given their beliefs) they cannot do so as a response to the *reasons* that support doing so. If something like agential possibility is the morally relevant sense of possibility in the dictum *ought implies can*, then it turns out that an agent's beliefs can affect the deontic status of the agent's actions by constraining which actions are possible for the agent.²⁷

3.2 What's Subjective Obligation?

We've seen that it's difficult to characterize the sense in which an agent's objective moral obligations are "insensitive" to the agent's beliefs--we will return to this puzzle shortly. But first we should note that it's challenging to characterize the sense in which an agent's *subjective* moral obligations are *sensitive* to the agent's beliefs. In fact, in "Subjective Rightness," Smith surveys many of the problems with four common ways of conceptualizing subjective rightness.²⁸

²⁷ Hicks (2022).

²⁸ Smith (2010), pp. 75-86.

For example, we might be tempted to say that an action is subjectively right when it's the action that, from the agent's perspective, seems most likely to be objectively right. But Smith points out sometimes the action that seems most likely to be objectively right is also extremely *risky*; so, this way of conceptualizing subjective rightness doesn't allow the riskiness of an action to affect its subjective rightness. Smith describes similarly worrisome objections to other first-pass conceptions of subjective rightness.

On Smith's view, a moral theory's principles should be *usable*, in the sense that an agent can derive from them a prescription for action if they want to. Objective moral principles are usable by agents with true beliefs about their situation, as well as by agents who hold false beliefs about their situation. However, objective moral principles are not usable by agents who *lack* beliefs about the descriptive features of their situation, nor by agents who have only *partial credences* concerning the descriptive features of their situation. Thus, Smith thinks that principles of objective rightness must be supplemented by principles of subjective rightness. Principles of subjective rightness form a hierarchy of decision-guides (relative to a principle of objective rightness); subjective principles that are higher in the hierarchy are usable by agents who have more true beliefs about the descriptive features of their situation, while agents who have less information must use principles of subjective rightness that fall lower in the hierarchy. In other words, descriptively uncertain agents can use principles of subjective rightness to approximate the satisfaction of principles of objective rightness.

Smith's account of subjective rightness describes how a moral theory can remain usable by descriptively uncertain agents. However, Smith's principles of subjective rightness cannot help

an agent who experiences distinctively *moral* uncertainty. An agent who is descriptively certain but morally uncertain can still use a principle of objective rightness to derive a prescription for action *if they are motivated to*; the problem is that such an agent's moral uncertainty could easily undermine their motivation to conform to that principle of objective rightness. (Imagine that it's a moral truth *that wrongdoers must always be punished*. Still, if one is uncertain about whether that's a moral truth, one will be disinclined to use that principle to guide one's actions; following such a principle when one is uncertain about it could feel too "morally risky" to a morally conscientious agent.²⁹)

The upshot of this discussion of Smith is that on her account of subjective rightness,³⁰ the subjective rightness of an action is never affected by an agent's moral beliefs and credences, only by an agent's descriptive partial credences. Thus, on Smith's view, the subjectively right action for an agent who holds false moral beliefs (or inaccurate moral credences) will often be an action that the agent *regards* as morally wrong. One might think this is a good feature of Smith's account of subjective rightness, because it rules out horrifying actions being counted as "subjectively right" simply because they stem from an agent's horrifying moral beliefs. However, one might also think that this feature of Smith's account is a weakness; given that most of us believe some moral falsehoods, it turns out that most of us are sometimes *subjectively* obligated to act in ways we cannot morally endorse. We might not want to rule out the possibility that *some* false moral beliefs can affect one's subjective moral obligations, particularly if we're interested in moral principles that can guide the actions of real-life, highly flawed moral agents.

²⁹ For discussions of the concept of moral risk (and "moral hedging" in response to risk), see Hicks (2019) and Weatherston (2014).

³⁰ As well as others' accounts; see Harman (2015).

3.3 The Incoherence of the Distinction

Stepping back, we can see two main challenges to the coherence of the objective/subjective moral obligation distinction.

First, there is something odd about characterizing objective obligation as what one ought to do “given the facts”; after all, some of “the facts” are facts *about one’s beliefs*, and so this characterization of objective moral obligation doesn’t by itself rule out the possibility that one’s beliefs affect one’s objective moral obligations. Moreover, we have independent reasons to think that an agent’s beliefs sometimes play a role in determining what an agent objectively ought to do; at the very least, an agent’s beliefs can determine whether the agent’s action falls under a morally salient act-type, and can also constrain what’s possible (in the morally relevant sense of “possible”) for the agent. Thus, it’s far from obvious that one’s objective moral obligations have nothing to do with one’s beliefs. And once we recognize that one’s beliefs can affect one’s objective moral obligations, it’s unclear whether we need a supplementary “subjective” conception of moral obligation.³¹

Second, notice that when it comes to subjective moral obligation, we face difficult choices about how to draw the line between those beliefs that affect one’s subjective moral obligations and those that don’t. In fact, we face many of the same line-drawing worries that motivated introducing the objective/subjective distinction in the first place.³² Thus, introducing a distinction between subjective and objective moral obligation doesn’t relieve us of the burden of determining which sorts of beliefs and mental states are “morally relevant” in cases like

³¹ Hicks (2022).

³² See pages 3-4 [CHANGE WHEN PAGINATION CHANGES] of this chapter.

Accidental Orchid and Intentional Orchid--we still have to make the same determinations at the level of *subjective* moral obligation.

At this point, the subjective/objective distinction does not appear to do the theoretical work philosophers originally wanted it to do. It doesn't provide a helpful analysis of our dueling intuitions about deontic status, or of our intuitions about blameworthiness. The distinction also doesn't solve the fundamental problem that a substantive moral theory cannot reliably guide the actions of highly imperfect moral agents. And more generally, the distinction doesn't prevent us from having to do the hard work of determining which types of ignorance are morally relevant. All of these observations suggest that our remaining work lies in first-order normative theorizing about the moral relevance of ignorance, *not* in metaethical theorizing about types of moral obligation.

4. New Directions

In light of the objective/subjective distinction's failure to do the theoretical work we hoped it would do, we ought to return to the original hard problem: to determine the conditions under which an agent's ignorance is morally relevant. To do this, we must engage in normative ethical theorizing that focuses on how beliefs and mental states affect the deontic statuses of actions. We have already seen two examples of how normative ethical theorizing can illuminate the role that an agent's beliefs play in determining what that agent morally ought to do. First, an agent's beliefs play an indispensable role in determining the *types* under which an agent's actions fall; thus, an objective theory of moral obligation that treats act-types as morally relevant is also a belief-sensitive moral theory. Second, beliefs (along with other mental states) arguably constrain

which actions are *possible* for an agent; given that an agent's beliefs affect what the agent can do, and given 'ought' implies 'can,' it follows that an agent's beliefs affect what that agent objectively ought to do.

No doubt there are other roles that beliefs might play in determining an agent's objective moral obligations. To uncover those roles, we might turn to the non-ideal theory tradition in moral and political philosophy.³³ Non-ideal moral theories are *objective* moral theories; non-ideal theorists do not usually take themselves to be theorizing about "secondary" subjective obligations.³⁴ At the same time, non-ideal theorists examine the ways in which the imperfections of agents--including various forms of ignorance--can shape those agents' moral obligations. Thus, the non-ideal theory tradition has already recognized that objective moral obligations can be shaped by the imperfections (including epistemic imperfections) of moral agents.

Instead of drawing a distinction between objective and subjective moral obligation, we can further explore the contours of a belief-sensitive *objective* moral theory. But one might object that this type of normative ethical theorizing won't help us with the problem of action-guidance--after all, the edicts issued by an objective, belief-sensitive moral theory will still be epistemically inaccessible to many moral agents. This is so for two reasons: (a) a moral theory that's belief-sensitive (that takes an agent's beliefs into account when assigning deontic statuses to that agent's available actions) is not necessarily a theory that simply directs the agent to "act on" their beliefs; (b) moreover, not all agents have reliable epistemic access to their own beliefs.

³³ Mills (2005); Rivera-López (2013); Tessman (2010).

³⁴ Berg (2018) is an exception.

This objection is correct--an objective, belief-sensitive moral theory is not reliably action-guiding. In fact, an objective, belief-sensitive moral theory can fail to guide action in all three ways described earlier in this chapter.³⁵ But recall that introducing a “subjective” conception of moral obligation didn’t make a moral theory much more action-guiding. Moreover, recall that there is a trade-off between action-guidance and substance; a substantive moral theory (one that does not simply direct agents to act on their beliefs) cannot be reliably action-guiding, and a reliably action-guiding theory cannot have much substance.³⁶ So, we might want to revisit our standards for action-guidance. Perhaps the idea that a moral theory must guide action in the sense that it must be “consultable” by people with mistaken beliefs fails to capture how high quality moral deliberation works. Perhaps the best way for someone to act as they ought is for that person to develop sensitivity to the moral reasons that are operative in their context, rather than for that person to consult a list of decision-guides.³⁷

One might also object that we ought not further explore belief-sensitive objective moral theories because they do not illuminate the nature of blameworthiness. But as we’ve seen, theories of subjective moral obligation also don’t illuminate the nature of blameworthiness, in part because there are many different conceptions of blame that need to be examined on their own terms (rather than being subsumed by the concept of subjective moral obligation). Moreover, even if we want a moral theory to offer an account of blameworthiness in addition to providing an explanation for why actions have the deontic statuses they have, it isn’t obvious that the theory’s account of blameworthiness must be *based* upon the theory’s account of deontic status; it’s

³⁵ See section 1.2.

³⁶ See section 2.2.

³⁷ Dancy (2006).

perfectly coherent for a moral theory's account of deontic status to be independent of the theory's account of culpability.

One might also object that this chapter has overlooked a significant motivation for introducing the objective/subjective moral obligation distinction: namely, that the distinction helps us make sense of how *advice* works. If we're working with an objective, belief-sensitive moral theory, then there will be cases in which an agent objectively ought to Φ , even though that agent would have a different obligation if they knew more.³⁸ But imagine that the ignorant agent who is about to Φ asks a better-informed person for advice. Sensibly, the better informed person might reply, "What you *really* ought to do is Ψ ." So, it looks like Ψ -ing is the agent's objective obligation, not Φ -ing, which suggests that an objective, belief-sensitive moral theory will get the extension of an agent's objective moral obligations wrong.

However, our intuitions about advice can be accounted for in a number of ways. First, we can say that the better informed person *makes it the case* that the agent's objective moral obligation is to Ψ ; in other words, we might think that the receipt of advice from better-informed people often *changes* one's objective moral obligations. Second, we can say that the intuition that the agent's *real* obligation all along was to Ψ is, in fact, an intuition that's tracking a counterfactual. It's certainly true (even before the agent seeks out advice) that *if the agent had known more then they would have been obligated to Ψ* . In light of the fickleness of our moral intuitions, the truth of that counterfactual claim could account for our intuition that the agent was objectively obligated to Ψ all along.

³⁸ Hicks (forthcoming).

Although this chapter has argued against drawing a distinction between objective and subjective moral obligation, we ought not rule out the possibility that there are similar, more coherent distinctions in the neighborhood. For example, it's sensible to distinguish between an agent's *moral* obligation and their *rational* obligation; and in fact, some philosophers who distinguish between objective and subjective moral obligation appear to treat subjective moral obligation as the same thing as (or coextensive with) rational obligation.³⁹ Perhaps there are cases in which an agent morally ought to Φ , and yet rationally (given their beliefs) ought to Ψ . This is a coherent possibility; but to capture its coherence, we need not introduce two distinct conceptions of *moral* obligation. Moreover, it's far less controversial to posit competing 'oughts' that operate in distinct normative domains (such as the moral 'ought' and the rational 'ought') than to posit competing 'oughts' that operate in the same normative domain (such as the subjective moral 'ought' and the objective moral 'ought'). It's plausible that many of the remaining intuitions that appear to support the objective/subjective moral obligation distinction in fact only support distinguishing between 'oughts' in distinct normative domains.

References

Berg, A. "Ideal Theory and 'Ought Implies Can'," *Pacific Philosophical Quarterly* 99, No. 4 (2018).

Dancy, J. *Ethics Without Principles* (Oxford University Press, 2006).

Feldman, F. "True and Useful: On the Structure of a Two Level Normative Theory," *Utilitas* 24, No. 2 (2012).

Graham, P. "In Defense of Objectivism about Moral Obligation," *Ethics* 121, No. 1 (2010).

³⁹ Lord (2015) argues that *deliberative* obligations are "perspectival" (subjective); Schwan (2017) similarly focuses on all-things-considered practical obligations; Muñoz & Spencer (2020) assume that the "subjective ought" is the 'ought' of rationality. It is a conceptually open question whether any of these 'oughts' are *moral* oughts.

Harman, E. "The Irrelevance of Moral Uncertainty," in *Oxford Studies in Metaethics* Volume 10, ed. Russ Shafer-Landau (Oxford University Press, 2015).

Harman, E. "Does moral ignorance exculpate?" *Ratio* 24, no. 4 (2011).

Hicks, A. "Dispensing with the Subjective Moral 'Ought'," *Oxford Studies in Normative Ethics* Volume 11, ed. Mark Timmons (Oxford University Press, 2022).

Hicks, A. "Moral Hedging and Responding to Reasons," *Pacific Philosophical Quarterly* 100, No. 3 (2019).

Hooker, B. *Ideal Code, Real World: A Rule-consequentialist Theory of Morality* (Oxford University Press, 2000).

Jackson, F. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," *Ethics* 101, No. 3 (1991).

Jones, Karen and François Schroeter (ed.). *The Many Moral Rationalisms* (Oxford University Press, 2018).

Kauppinen, Antti, "Moral Sentimentalism," *The Stanford Encyclopedia of Philosophy* (Winter 2018 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2018/entries/moral-sentimentalism/>.

Lord, E. "Acting for the Right Reasons, Ability, and Obligation," in *Oxford Studies in Metaethics* Volume 10, ed. Russ Shafer-Landau (Oxford University Press, 2015).

MacAskill, W., Krister Bykvist, and Toby Ord. *Moral Uncertainty* (Oxford University Press, 2020).

Mill, J.S. *Utilitarianism* (Batoche Books, 1863).

Mills, C. "'Ideal Theory' as Ideology," *Hypatia* 21, No. 3 (2005).

Muñoz, D. and Jack Spencer. "Knowledge of Objective 'Oughts': Monotonicity and the New Miners Puzzle," *Philosophy and Phenomenological Research* (published online in 2020).

Olsen, K. "A Defense of the Subjective/Objective Moral Ought Distinction," *Journal of Ethics* 21 (2017).

Railton, P. "Alienation, Consequentialism, and the Demands of Morality," *Philosophy and Public Affairs* 13 (1984).

Regan, D. *Utilitarianism and Cooperation* (Oxford University Press, 1980).

- Rivera-López, E. "Ideal and Nonideal Ethics and Political Philosophy," in *International Encyclopedia of Ethics*, ed. Hugh LaFollette (Blackwell, 2013).
- Rosen, G. "Skepticism about moral responsibility," *Philosophical Perspectives* 18, No. 1 (2004).
- Schwan, B. "What ability can do," *Philosophical Studies* 175, No. 3 (2018).
- Shoemaker, D. "Attributability, Answerability, and Accountability: Toward a Wider Theory of Moral Responsibility," *Ethics* 121, No. 3 (2011).
- Sliwa, P. "The Power of Excuses," *Philosophy and Public Affairs* 47, No. 1 (2019).
- Smart, J.J.C. "Extreme and Restricted Utilitarianism," *The Philosophical Quarterly* 6, No. 25 (Oct. 1956).
- Smith, A. "Responsibility as Answerability," *Inquiry* 58, No. 2 (2015).
- Smith, H.M. *Making Morality Work* (Oxford University Press, 2018).
- Smith, H.M. "Subjective Rightness," *Social Philosophy and Policy* 27, No. 2 (2010).
- Tessman, L. "Idealizing Morality," *Hypatia* 25, No. 4 (2010).
- Väyrynen, Pekka, "Thick Ethical Concepts," *The Stanford Encyclopedia of Philosophy* (Spring 2021 Edition), ed. Edward N. Zalta, URL = <https://plato.stanford.edu/archives/spr2021/entries/thick-ethical-concepts/>.
- Weatherson, B. *Normative Externalism* (Oxford University Press, 2019).
- Weatherson, B. "Running Risks Morally," *Philosophical Studies* 167, No. 1 (2014).
- Zimmerman, M. "Is Moral Obligation Objective or Subjective?" *Utilitas* 18, No. 4 (2006).
- Zimmerman, M. "Moral responsibility and ignorance" *Ethics* 107, No. 3 (1997).