

1 Subjectivism, Objectivism, and Ecumenism

The debate between Subjectivists and Objectivists about moral wrongness is a debate about what kinds of facts the moral status of an agent's action – whether it is morally wrong or not – is true in virtue of, or is grounded in.¹ The terms “Subjectivism” and “Objectivism,” however, need precisifying.

As the Subjectivism/Objectivism debate is a debate about what facts the moral status of an action is grounded in, the first question to ask is: What are the different sets of facts participants to the debate are in dispute about? Subjectivists think that the relevant facts are those that constitute the agent's subjective circumstances, whereas Objectivists think the relevant facts are those that constitute the agent's objective circumstances. Roughly speaking, one's subjective circumstances just are all the facts concerning how the world seems to be from one's perspective. And, roughly speaking, one's objective circumstances just are all the facts concerning how the world actually is, independent of how things seem to be from one's perspective. As I'm understanding it, then, the set of facts that constitute a person's objective circumstances is the complement of that which constitutes their subjective circumstances – if a certain fact is among an agent's subjective circumstances, then it is not among their objective circumstances, and if a certain fact is not among an agent's subjective circumstances, then it is among their objective circumstances.²

Subjectivisms can vary depending on the way in which “how the world seems to be from the agent's perspective” is interpreted. *Belief Subjectivists* maintain that an agent's subjective circumstances comprise all the facts there are about what they believe about the world around them. *Evidence Subjectivists* contend that an agent's subjective circumstances comprise all the facts there are about what their evidence about the world around them is. For much of what follows, the distinction between Belief Subjectivism and Evidence Subjectivism won't much matter, but it is important to keep in mind that different Subjectivists take

¹ The notion of grounding is supposed to capture this “in virtue of” thought. A grounding relation is thus supposed to be a one-way asymmetric necessary determination relation stronger than, though entailing, supervenience. Modern classics on the nature and logic of this grounding relation are Rosen (2010) and Fine (2012).

² What about facts that are conjunctions and disjunctions of facts both in an agent's subjective circumstances and their objective circumstances? For example, take either the conjunction [(*S* believes that snow is white) & (snow is white)] or the disjunction [(*S*'s evidence indicates that snow is white) ∨ (snow is white)]. How should we classify these propositions? Though this is an interesting question, it isn't one that makes any important difference for the arguments I shall be considering. For this reason, we could, by stipulation, treat both conjunctions and disjunctions of combinations of subjective and objective facts as subjective facts. (The issue here is clearly related to those that surround A. N. Prior's [1960] famous demonstration that an “ought”-claim can indeed be derived from an “is”-claim.)

different stances on the kinds of subjective circumstances they take to be the ground of moral wrongness.³

The easiest way to get a grip on what is in dispute between Objectivists and Subjectivists about moral wrongness is to consider a case about which they disagree:

Switch: Gomez is looking for her wallet. Chen stands near the light switch and can flip it, thereby turning on the light and helping Gomez in her search for her wallet. Chen believes, and all of her evidence indicates, that were she to flip the light switch, all that would happen is that the light would go on and Gomez would be helped in her search for her wallet. In actual fact, the light switch has been hooked up to a bomb which will go off, killing twenty innocent people, if Chen flips it.

Subjectivists generally maintain that it would not be morally wrong for Chen to flip the switch in *Switch* – from Chen’s perspective, flipping the switch will hurt no one and help Gomez find her wallet. Objectivists, on the other hand, generally maintain that it would be morally wrong for Chen to flip the switch in *Switch* – flipping the switch will cause twenty innocent people to be killed and it is morally wrong to kill innocent people when one can avoid doing so.

Having gotten a flavor of what’s in dispute in the Subjectivism/Objectivism debate, I’ll now more carefully consider each of the possible positions in this debate one might take. Subjectivists maintain that the moral status of an agent’s action is grounded in their subjective circumstances. But because a fact can be either fully or partially grounded in some other set of facts, we need first to distinguish between two theses that might legitimately claim the mantle “Subjectivism”:

Subjectivism: the moral status of an agent’s action is *fully grounded* in their subjective circumstances.

Ecumenism: the moral status of an agent’s action is *merely partially grounded* in their subjective circumstances.

³ There is an in-depth discussion of the various different kinds of Subjectivism in Driver (2012). Some of those I call “Subjectivists” others sometimes call “Perspectivists” (Kiesewetter 2011, 2017; Lord 2015). Others sometimes use “Subjectivism” to pick out what I do by “Belief Subjectivism” (Zimmerman 2008, 2014). There is some reason to use “Perspectivism” to pick out what I pick out by “Subjectivism” just because there are a number of different views in ethics, and philosophy more generally, that might well be called “Subjectivist.” (According to some views, a moral theory counts as “Subjectivist” just in case according to it the moral status of an action is partially grounded in motivational features of the agent at the time of its performance. That is not the sense of “Subjectivism” I am employing.) However, I stick with the “Subjectivism/Objectivism” terminology because I think it connects up more broadly with views that have been called versions of “Subjectivism” throughout the literature on this particular debate about moral wrongness.

To say that the moral status of an agent's action is fully grounded in their subjective circumstances is to say that every one of the facts that altogether ground the moral status of the agent's action are facts in their subjective circumstances. To say that the moral status of an agent's action is merely partially grounded in their subjective circumstances is to say that some of the facts that altogether ground the moral status of the agent's action are facts in their subjective circumstances, while some of the facts that altogether ground the status of the agent's action are facts that are not in their subjective circumstances. (Clearly, given how we've defined an agent's subjective and objective circumstances, Ecumenists hold that the moral status of an agent's action is merely partially grounded in their objective circumstances as well.) Though some might claim that endorsing anything as strong as Ecumenism is enough to count one as a Subjectivist, for the purposes of our discussion only those who endorse Subjectivism, along the lines defined above, according to which even merely partial grounding in an agent's objective circumstances is ruled out, will count as Subjectivists.

The definition of Objectivism follows naturally:

Objectivism: the moral status of an agent's action is *fully grounded* in their objective circumstances.

Just as our initial definition of "Subjectivism" has it that the moral status of an agent's actions is fully grounded in their subjective circumstances, this definition of "Objectivism," correspondingly, has it that the moral status of an agent's actions is fully grounded in their objective circumstances.

So understood, Ecumenism and Objectivism are both incompatible with Subjectivism and with each other. Our three positions are mutually exclusive and mutually exhaustive. Anything other than full grounding of the moral status of an agent's action in their subjective circumstances is incompatible with Subjectivism. No grounding, either full or merely partial, of the moral status of an agent's action in anything other than their objective circumstances is compatible with Objectivism. And only merely partial grounding of the moral status of an agent's action in their subjective circumstances is necessary and sufficient for the truth of Ecumenism. In what follows, I shall offer a qualified defense of Objectivism. In particular, I'll consider various arguments first against Subjectivism and then against Objectivism. Though there are arguments that do seem to establish the falsity of Subjectivism, none of the arguments against Objectivism definitively show it to be false. Now, for all that will be said, it might well be that one of the arguments against Objectivism does in fact go through, in which case it would be some form of Ecumenism that is true. As it is unclear whether any of those arguments do succeed, all that will be offered is merely a qualified defense of Objectivism.

The debate between Subjectivists, Objectivists, and Ecumenists is of fundamental significance.⁴ Whether what I may permissibly do is grounded in how things seem to me or in how the world actually is, or in both, is of paramount importance for moral theory. If Subjectivism is true, then a whole swath of moral theories, including, for instance, all forms of Objective Consequentialism – the view according to which any option one has that fails to maximize value, of all the options one has, is morally wrong – are ruled out simply as a matter of course. It should be a top priority, then, to settle this question before embarking on an in-depth investigation into the grounds of moral wrongness, so that we don't go wrong in the very first step of our investigation.⁵

2 The Sense-Splitting Reply

Some don't see the Subjectivism/Objectivism debate as being of deep moral significance. They are inclined to think that trying to settle this dispute is a fool's errand because the debate between Subjectivists and Objectivists about moral wrongness is in fact much ado about nothing. They diagnose the debate as grounded in a confusion. The confusion is about wrongness itself – in particular, the confusion involves thinking that there is a single notion of wrongness about which Subjectivists and Objectivists are disagreeing. Rather, so goes this reply, Subjectivists and Objectivists are just talking past each other because they are each concerned with a distinct kind of wrongness. There is a subjective notion of wrongness and an objective notion of wrongness, and each side of the debate is correct about their respective notions of wrongness.⁶

This sense-splitting reply misses its mark. Whether there are subjective and objective notions of wrongness or not, we can pin down the notion of wrongness about which Subjectivists and Objectivists are in dispute, and it is clear, once it is

⁴ Examples of Subjectivists, in the sense I've argued for here, are Prichard (2002), Ross (1939), and Howard-Snyder (2005). Objectivists include Moore (1912), Ross (2002), Feldman (1986), Thomson (1990), and Graham (2010). Jackson (1991) is an Ecumenist because he thinks that moral wrongness is grounded both in the agent's beliefs about nonmoral matters and in the objective values of the states of affairs the agent believes to obtain. Zimmerman (2008, 2014) and Kiesewetter (2011, 2017) are predominantly subjective Ecumenists because, though they think that the moral wrongness of an action is primarily a function of the agent's evidence, they think it is also a function of the options the agent has available to them, and what those options are is an objective fact. Zimmerman (2008, 2014) uses the term "Prospectivism" to pick out the kind of predominantly subjective Ecumenism he favors. Another proponent of Prospectivism is Mason (2013).

⁵ Similar issues to those that divide Subjectivists, Objectivists, and Ecumenists also arise within the literature on normative reasons and normative "oughts" (both practical and epistemic), including Sepielli (2012, 2018), Kiesewetter (2011), Lockhart (2000), Lord (2017, 2018), and Way and Whiting (2016, 2017).

⁶ This sense-splitting strategy is employed to in Sepielli (2012).

so pinned down, that they are not talking past each other. What, then, is the notion of wrongness about which Subjectivists and Objectivists are in dispute? It is that of which we are trying to give an account when doing moral theory (i.e., it is the notion of wrongness about which Kantians and Utilitarians disagree when they give their respective accounts of moral wrongness). And that notion of wrongness is the notion of wrongness that is of ultimate concern to the morally conscientious person when in their deliberations about what to do they ask themselves, “What would be morally wrong for me to do in this situation?”⁷ The morally conscientious person, who asks themselves this question in their deliberations about what to do, isn’t actually asking themselves two distinct questions. Nor are they muddle-headedly thinking in terms of some confused amalgam of two distinct questions. No. There is one unambiguous question they are asking themselves. And, thus, there is only one notion of wrongness they are ultimately concerned about when in their deliberations about what to do they want to avoid doing something morally wrong. It is *that* notion of wrongness about which Subjectivists and Objectivists are in dispute. In other, more explicit words, Subjectivists think that the notion of wrongness that is of ultimate concern to the morally conscientious person when, in their deliberations about what to do, they ask themselves, “What would be morally wrong for me to do?”, is one that is grounded solely in facts in their subjective circumstances. Objectivists, on the other hand, think that the notion of wrongness that the morally conscientious agent is asking about when they ask themselves that question is one that is not grounded in any facts in their subjective circumstances. And Ecumenists think that the kind of wrongness with which the morally conscientious agent is ultimately concerned in their deliberations about what to do is grounded in a combination of their subjective and objective circumstances. As we can independently anchor the one and only notion of wrongness about which Subjectivists, Objectivists, and Ecumenists are in dispute to the notion of wrongness about which the morally conscientious agent is ultimately concerned in their deliberations about what to do, the sense-splitting reply to the Subjectivism/Objectivism debate collapses. There *is* a unique notion of wrongness about which the parties are disagreeing and theirs is very much a live dispute.⁸

⁷ It is important to distinguish the morally conscientious person from the morally virtuous person, or the moral person *simpliciter*. A morally conscientious person is one who is particularly concerned, punctiliously so, about moral wrongness. They are concerned to avoid wrongdoing. They needn’t be concerned about anything like maximizing value (unless they happen to be a Consequentialist). Nor need they be at all interested in going above and beyond the call of moral duty; doing the morally supererogatory isn’t something one need be concerned about to count as morally conscientious. Zimmerman (2008) is the first to note that it is the deliberations of the morally conscientious agent in particular that are central to the Subjectivism/Objectivism debate.

⁸ This response to the sense-splitting reply is employed in Zimmerman (2008). A version of this reply is also to be found in Kiesewetter (2011, 2017) and Lord (2015, 2018).

3 Against Objectivism

There are a number of different arguments against strict Objectivism about moral wrongness. In this section, I survey and evaluate the prospects of various of these arguments. My conclusion is that none of these arguments against Objectivism is conclusive.

3.1 The Brute Intuition Argument against Objectivism

One of the most popular arguments against Objectivism involves a brute appeal to intuition: it is not morally wrong for Chen to flip the light switch in *Switch*, but, as it would be morally wrong for Chen to flip it in that case if Objectivism were true, Objectivism is false. This argument is simple and straightforward. It is also, however, not very dialectically effective. Objectivists simply reject the intuition that Chen does not act morally wrongly in flipping the switch in *Switch*. True, Chen acts blamelessly in doing so, but that's just an instance of blameless wrongdoing. And according to Objectivists, blameless wrongdoing is ubiquitous. For instance, for the Objectivist, when I walk off with your jacket, one that is identical to my own, by mistake from the cloakroom, I wrongly violate your right against me that I not interfere with your property. But because there was no way I could have known that I was taking your jacket (suppose that that is the case in this scenario), I'm not to blame for taking it. In such a case, indignation toward me on the part of others on account of my behaving the way I did would be totally out of place.⁹ Chen's flipping the switch in *Switch* is, for the Objectivist, just another instance, though one with more disastrous consequences, of this more general kind of blameless wrongdoing.¹⁰

Not only can there be blameless wrongdoing, the Objectivist will say, there can even be cases in which one would be blameworthy for failing to do that which it is morally wrong for one to do. Chen, for instance, not only would not be blameworthy for flipping the switch in *Switch*, but she would, in fact, be blameworthy for not flipping it. Given that she believes, and all of her evidence indicates, that flipping the switch will do nothing but help Gomez, and, in

⁹ In arguing on the Objectivist's behalf here, I'm implicitly appealing to something like a Strawsonian conception of moral blameworthiness (Strawson 1962), one according to which a person is blameworthy for doing something if and only if it would be appropriate for one of the blame emotions – guilt, resentment, and indignation – to be borne toward them on account of their doing that thing. For recent developments of this approach to moral blameworthiness see Wallace (1994), Graham (2014), and Rosen (2015). Though I appeal to this conception of moral blameworthiness here, a similar point could be made, I think, with respect to other non-emotion-centered conceptions of moral blameworthiness, like those of Zimmerman (1988) and Scanlon (2008).

¹⁰ The Objectivist maneuver of distinguishing wrongdoing and blameworthiness is most famously deployed by Moore (1912).

general, she morally ought to help Gomez if she can do so at no cost to her or anyone else, it would be morally blameworthy for Chen to refrain from flipping the switch, even though doing so would, according to the Objectivist, be morally wrong. For the Objectivist, then, moral wrongness and moral blameworthiness can, and often do, come apart drastically. Because of this, the argument against Objectivism consisting in the brute appeal to the intuition that it would not be morally wrong for Chen to flip the switch in *Switch* is particularly dialectically ineffective.

3.2 The Action-Guiding/Usability Argument against Objectivism

Another argument against Objectivism appeals to the thought that to be true a moral theory must be action-guiding.¹¹ The thought is that being inherently practical, as opposed to merely theoretical, morality, of necessity, must provide agents with guidance about what to do in any situation in which they might find themselves. Because the truth of Objectivism about moral wrongness would render morality incapable of carrying out this function, so goes this argument, Objectivism is false. (It is important to this argument that morality must necessarily be *universally* action-guiding, in that it must provide guidance in all cases, including ones in which agents don't have full knowledge of their situation. This is because in cases of full knowledge, presumably, objective moral theories are no less action-guiding than are subjective ones.) The correct moral theory must be *usable* in helping one determine what to do. Knowing the principles that determine moral rightness and wrongness cannot leave one hopelessly helpless in the face of ignorance and uncertainty. But that's precisely the position one would often find oneself in were Objectivism true. So, Objectivism is false.

This argument may be persuasive from a distance. But upon closer inspection its weaknesses are manifest. First, though it is affirmed both that to be true a moral theory must be action-guiding and that Objectivist moral theories are clearly not action-guiding, usually very little is said either about what it means for a moral theory to be action-guiding or about why it is that Objectivist moral theories can't be action-guiding in that sense.¹² What exactly does it mean to say that a moral theory must be usable for it to be true? Usable how? And by whom? Unless these questions can be satisfactorily answered, this argument against Objectivism is at best a promissory note, and at worst a mirage.

¹¹ A recent attempt at this style of argument is Fox (2019).

¹² A notable exception is Smith (2018). Smith goes into very particular detail about what it means to say that a moral theory is usable and why it is that purely objective moral theories fail to be usable in the relevant sense. Though Smith is clear about what it takes for a moral theory to be usable, it is less clear whether she establishes that morality must indeed be usable in that sense for it to be true.

Take, for example, the question: by whom must a moral theory be usable for it to be usable in the sense necessary for its being true? Consider a very small child who doesn't possess the concepts of moral rightness and wrongness. It is dubious whether, for any sense of "usable," *any* moral theory could be usable by such a child. Presumably, then, a moral theory needn't be usable by anyone for it to be true. So, maybe to be true a moral theory need only be usable by those who possess the concepts employed in the principles that constitute it.¹³ Arguably, morality only applies to those who possess the concepts of rightness and wrongness; that is, the only beings who can act morally rightly or wrongly are ones who possess those concepts, and so, perhaps, because of that, those for whom a moral theory must be usable for it to be true are all and only those who possess the concepts of rightness and wrongness.¹⁴ But consider another person – someone who, though they possess all the relevant concepts, invariably makes logical mistakes whenever they derive particular consequences from general principles. Arguably, for such a person no moral theory would be usable, for any sense of "usable." Okay. So maybe the usability requirement is meant to exclude such agents; perhaps a moral theory must be usable by a minimally idealized rational agent who possesses the concepts employed in it. That may be a defensible restriction, but what exactly is the relevant idealization and is it consistent with the motivation, whatever it happens to be, for the thought that moral theories must be usable, or action-guiding, for them to be true? This is not so easy a question to answer, for the motivation for the usability criterion is often unspecified.

Consider also the question *how* a theory must be usable for it to be true. What does it mean for a moral theory's principles to be usable by an agent? What counts as a person's using a moral theory to determine what to do? Must it be the case that for any situation in which one might find oneself one is able to determine or deduce from the theory what to do in that situation? Presumably not, for in many cases there will be no action that the moral theory dictates that the agent must perform. Take, for instance, my choice of having either orange juice or apple juice with my breakfast this morning. Many moral theories surely are no help to me in deciding this question. Does that make them not usable? That would be implausible. Maybe it need only be the case that a theory be such that one can use it to decide what to do in cases in which the theory requires that one do some particular thing. The choice between orange and apple juice, being

¹³ This presumably requires more than just the possession of the concepts of rightness and wrongness, for a person who possessed those concepts but who lacked the non-deontic concepts in which the theory declared them to be grounded could just as well not use that theory as if they lacked the concepts of rightness and wrongness themselves.

¹⁴ That morality only applies to those who possess the concepts of rightness and wrongness is often thought to explain the widely held belief that it is conceptually impossible for very small children and most nonhuman animals to act rightly or wrongly.

one that no plausible moral theory would take a stand on, isn't a choice situation in which a moral theory needs to be usable in determining what to do for it to count as usable in the relevant sense. But a choice between doing that which would maximize utility and that which would keep a promise very plausibly is one that a plausible moral theory will take a stand on; that is, it is one for which the moral theory will dictate that one is required to do one or required to do the other. But consider a theory according to which, because the relevant moral factors balance out, both keeping the promise and maximizing utility were deemed permissible. Would such a moral theory be unusable in this situation? That is, would it be the case that because one couldn't figure out what to do in the situation from the moral theory that the theory is not universally usable? Again, that would be implausible.

Crucially, for the proponent of the usability argument against Objectivism, moral theories, to be true, must be usable in situations of uncertainty. But it's not so clear that a purely objective moral theory can't be usable in situations of uncertainty. Suppose Lopez faces a choice between pressing a certain button and doing nothing and she knows that pressing the button will have a 50 percent chance of producing more utility than her doing nothing and a 50 percent chance of doing nothing. In such a case, it seems Lopez most certainly *can* use Objective Utilitarianism – according to which what she morally ought to do is maximize utility – to determine what to do in this situation; viz., press the button. In fact, if Lopez is uncertain what to do because she is uncertain which moral theory is true, then comes to believe Objective Utilitarianism, and, then, because of that belief, is moved to press the button, then surely she has used Objective Utilitarianism in deciding what to do. But if she's used the theory to decide what to do in her situation, then the theory is usable for her in deciding what to do in that situation involving uncertainty.

But even if an objective moral theory can be usable in some situations of uncertainty, does it follow that an objective moral theory can be usable in *all possible* situations of uncertainty? Well, again, it all depends on how one understands what it is to use a moral theory to determine what to do. On one understanding of that notion, arguably, objective moral theories are indeed usable. Begin with an analogy. Take the purely self-interested person – one who is solely concerned with their own well-being – facing a choice under conditions of uncertainty about the outcomes of their various options. Suppose this person doesn't have a settled view as to what the correct theory of well-being – a theory of that in virtue of which a person is as well off as they are – is. Then suppose they come to accept hedonism – the theory according to which how well off one is is a straightforward function of the pleasures and pains in one's life – as the correct theory of well-being. They then rationally use the various evaluations of their well-being in the various possible

outcomes of each of their options as inputs into some reasonable decision theory, and choose that option which is the output of that decision theory. In this case, it surely seems as if the purely self-interested person has in fact used hedonism to help them determine what to do in their situation of uncertainty. But, and here's the point, hedonism is a purely objective theory of well-being – it doesn't tie an agent's well-being either to their beliefs about, or their evidence concerning, the world around them. So, purely objective theories of well-being can be used in situations of uncertainty to help purely self-interested persons decide what to do, and thus purely objective theories of well-being can be usable in such situations. Now, why should the case of the morally conscientious person and objective moral theories be any different? If a morally conscientious person takes whatever objective moral theory they believe to be true and inputs its verdicts about the various possible outcomes of their different options into a reasonable decision theory, and then chooses that option which is the output of that theory, then how have they not used that purely objective moral theory to help them determine what to do?

The proponent of the usability argument against Objectivism is likely to protest here that in the situation as described, the morally conscientious person has not *solely* used the purely objective moral theory to determine what to do, but rather has instead used the theory *and the reasonable decision theory* together to determine what to do. Maybe, they might maintain, for a moral theory to be usable in a situation it must be the case that an agent can use the theory, *and the theory alone*, to determine what to do in their situation. But this can't really be the criterion, for even subjective moral theories are not such that one can use them *alone* to determine what to do in one's situation. Even subjective moral theories have to be coupled with deductive reasoning – modus ponens, modus tollens, etc. – for anyone to use them to determine what to do. But if, to be usable in the sense necessary to be true, according to the usability argument against Objectivism, a moral theory can be used in conjunction with deductive reasoning to determine for an agent what to do in any situation in which they find themselves, then why too couldn't a moral theory count as usable in the sense required for it to be true if it can, in conjunction with reasoning in accord with a reasonable decision theory, determine for an agent what to do in any situation in which they find themselves? In other words, if deductive reasoning is allowed in, then why can't reasoning in accord with a reasonable decision theory be let in too? The proponent of the usability objection to Objectivism hasn't yet motivated an account of usability that is both plausible and which rules out objective moral theories as unusable in situations of uncertainty.¹⁵

¹⁵ One common response to the Action-Guiding/Usability argument against Objectivism is to endorse a “two-level” moral theory. This approach is favored by Feldman (2012) and Smith (2018). As Feldman puts it, a two-level moral theory is one that incorporates into it both a “theoretical-level