




ARTICLES

From Contextualism to Contrastivism in Moral Theory

Walter Sinnott-Armstrong 

Department of Philosophy, Duke University, Durham, NC, USA
Email: ws66@duke.edu

Abstract

In *Morality by Degrees*, Alastair Norcross presents contextualist accounts of good and right acts as well as harm and free will. All of his analyses compare what is assessed with “the appropriate alternative,” which is supposed to vary with context. This paper clarifies Norcross’s approach, distinguishes it from previous versions of moral contextualism and contrastivism, and reveals difficulties in adequately specifying the context and the appropriate alternative. It also shows how these difficulties can be avoided by moving from contextualism to a kind of contrastivism that does not claim that any alternative is or is not appropriate or relevant.

Keywords: contextualism; contrastivism; scalar consequentialism; moral theory

Consequentialisms have been breeding like rabbits in recent years (Portmore 2020). Many of these new versions of consequentialism attempt to solve problems for earlier formulations. One prominent kind of challenge invokes purported counterexamples where consequentialism seems to imply moral judgments that seem obviously incorrect. Many consequentialists respond to such challenges by either (a) biting the bullet (accepting the counterintuitive implications while perhaps also trying to make them more palatable) or (b) consequentializing (modifying standard consequentialism so that it can accommodate commonsense moral judgments) (Sinnott-Armstrong 2023).

A more radical approach is adopted by Alastair Norcross (2020) in his small but grand book, *Morality by Degrees: Reasons without Demands*. (All parenthetical references are to this book unless otherwise noted.) There Norcross develops a “scalar consequentialism” (11) in normative ethics and “contextualism” (108ff) in moral semantics. This combination is supposed to avoid both horns of the dilemma – (a) and (b) – by changing the question.

According to Norcross, traditional consequentialists and their critics get into trouble by asking questions out of context. They focus on simple questions of whether acts are right or wrong, good or bad, and beneficial or harmful without regard to who asks those questions under which circumstances. Norcross rejects these decontextualized questions because the context is necessary to determine which alternatives are being contrasted with the assessed act. In Chapter 2, Norcross explains why “consequentialism should

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

not employ the notions of right and wrong” (48; cf. 28–35, 40–47). Chapter 3 then argues that “consequentialism cannot provide a satisfactory account of the goodness of actions” and “strictly speaking, a consequentialist cannot judge one action to be better or worse than another action performed at a different time or by a different person” (82). Chapter 4 makes related points about the notion of harm and concludes, “There is no fundamental non-comparative moral fact of the form ‘act A harms person X’” (101). If these bold claims are correct, then almost everyone engaged in moral philosophy – not only consequentialists – will need to rethink their entire enterprise.

A paradox

As an example, let us consider how Norcross argues against the popular view that “An act A is good iff the world would have been worse if A hadn’t been performed” (112). Norcross asks, “What would the world have been like, if Agent hadn’t [done A]?” and answers, “That depends on what Agent would have done instead” (112). To illustrate the problem, he discusses this case:

Burning Building: There are ten people trapped in a burning building. Agent can rescue them one at a time. Each trip into the building to rescue one person involves a considerable amount of effort, risk and unpleasantness. It is possible, albeit difficult and risky, for Agent to rescue all ten. Suppose Agent rescues nine people and then stops, exhausted and burned. She could have rescued the tenth, so she doesn’t do the very best she can, but do we really want to say that her rescue of nine people wasn’t good (was actually bad)? (113)

Norcross assumes that his readers (like him) will find it counterintuitive to judge that Agent’s rescue of nine people was simply not good (or was simply bad). The problem is that it also seems counterintuitive to judge that Agent’s rescue of *only* nine people was simply good, since Agent could have rescued all ten. Neither simple (non-contextual) judgment seems accurate, so we seem to be stuck in a paradox.

Critics might object that rescuing nine people and rescuing *only* nine people are different acts, so there is no paradox if the former is good and the latter is not good (cf. 115). Admittedly, rescuing nine is a different act than stopping after rescuing nine, because the stopping occurs at a later time than the rescuing, so the rescuing and the stopping can be evaluated differently. However, the paradox still arises for rescuing nine versus rescuing *only* nine, at least if those descriptions refer to the same action. Some fine-grained theories of act individuation (Goldman 1970; Kim 1976) deny that rescuing nine people and rescuing *only* nine people are the same act, partly because Agent could rescue (at least) nine without rescuing *only* nine simply by rescuing ten. However, it would be very misleading to say that Agent performed one act of rescuing nine and also performed another act of rescuing *only* nine. Agent’s act of rescuing the first person is distinct from Agent’s act of rescuing the second person, and so on; but the people who were rescued when Agent rescued nine and those rescued when Agent rescued *only* nine are the same people. Agent did not rescue eighteen people. Moreover, Agent did not perform two separate rescues when she rescued nine and also rescued *only* nine. The rescued people were rescued only once. If a rescued person owes a reward of \$1000 to their rescuer every time they are rescued, then they owe Agent \$1000 for rescuing them on that one occasion – not \$2000 for two acts of rescuing.

All of this suggests that Agent performed only one act at the time when she rescued nine and *only* nine people. That suggestion is supported by theories that individuate acts by their spatio-temporal locations (Anscombe 1957; Davidson 1985). If Agent rescues nine and *only* nine people, then the act of rescuing nine people and the act of rescuing *only* nine people occur at exactly the same time and place, so this is a single act of rescuing, even if that act can be described in various ways.

This conclusion restores the paradox. It seems counterintuitive to judge that Agent's rescue of nine and *only* nine people was simply not good, but it also seems counterintuitive to judge that Agent's rescue of nine and *only* nine people was simply good. Similar arguments lead to similar paradoxes regarding which acts are harmful (Chapter 4, 115–16) and which acts are wrong (Chapter 2), especially for consequentialists who judge acts to be right when they are good and wrong when they are bad.

A way out of the paradox

Norcross's solution appeals to two claims that are central to his contextualism (though these two claims can also be accepted by non-contextualists, as we will see):

Contrast Claim: Any unified theory [of good action] requires the fixing of a contrast point. (113)

Comparison Claim: our (moral) reasons for choosing between alternative actions, institutions, etc. are essentially comparative and correspond to the comparative consequential value of the options. (108)

Although neither claim is perfectly perspicuous, I will interpret the contrast claim as saying that acts are never good (or not) simply by themselves but are always good (or not) only in contrast with certain alternatives – and similarly for bad, wrong, harmful, etc. I will also focus on the part of the comparison claim that says that alternatives, including those required by the contrast claim, should be compared on the basis of “the comparative consequential value of the options.”

So understood, these claims are independent. One could consistently accept the contrast claim but not the comparison claim if one required alternatives but sometimes compared them in some way apart from their consequential value (Snedegar 2017). In the other direction, one could consistently accept the comparison claim but not the contrast claim if one allowed some judgments without contrasts but also allowed other judgments with contrasts, and then always compared the alternatives by their consequential value. Despite their independence, Norcross endorses both claims.

Together these two claims dissolve the paradox. Neither the judgment that Agent's rescue of nine and only nine was simply good nor the judgment that this act was simply not good (or was simply bad) includes a “contrast point” as required by the contrast claim. The only kinds of judgments that would be allowed by the contrast claim would be judgments like the judgments that Agent's rescue of nine was good *in contrast with the alternative of rescuing eight or fewer* and was not good (and was bad) *in contrast with the alternative of rescuing ten*. The comparison claim then implies that both of these contrastive judgments are true, because the consequential value of rescuing nine is worse than that of rescuing ten but better than that of rescuing eight or fewer. These resulting contrastive judgments do not even seem counterintuitive, so the paradox disappears.

A new problem

The trick in dissolving the initial paradox was to avoid non-contrastive judgments, which are judgments that an act is or is not good (or bad, wrong, harmful, etc.) that do not specify any contrasting alternative. This side step, however, introduces another problem. Everyday conversations typically include non-contrastive judgments and only rarely explicitly specify contrasting alternatives. Most people say things like “This act was good” instead of “This act was good in contrast with this alternative” (and similarly for bad, wrong, harmful, etc.). The new problem is to specify how we should understand such non-contrastive judgments.

One possibility is to hold that non-contrastive judgments are never true, either because they are all false or because they are all neither true nor false, perhaps because they are unclear or nonsensical (109). This *error theory* could be used to support *eliminativism*, which claims that we should eliminate or avoid all non-contrastive judgments (although it might be useful to keep making non-contrastive judgments even if none is true). In any case, Norcross expresses some sympathy with error theory but then rejects it along with eliminativism in everyday discourse, although he does suggest that we should eliminate non-contrastive judgments from the fundamental level of consequentialist moral theory (110).

Instead of error theory and eliminativism, Norcross suggests “a reductivist account of these notions, from which it follows that it is possible, even quite common, to express substantively true or false propositions involving them” (110). Although it is not clear exactly what reduces to what, Norcross’s basic idea is that non-contrastive sentences at the non-fundamental level of everyday discourse can be used to express true or false contrastive propositions at the fundamental level of moral theory (118, 124). Compare an everyday non-contrastive judgment that “This child is tall” being used to express a true or false contrastive proposition that this child is tall in contrast with other people her age, which does not imply that this child is tall in contrast with adults or that this child is tall without qualification. Analogously, Norcross seems to hold that an everyday non-contrastive sentence like “Agent’s rescue of nine was good” can be used to express a true contrastive proposition that Agent’s rescue of nine was good in contrast with rescuing eight or fewer (or, in a different context, a false contrastive proposition that Agent’s rescue of nine was good in contrast with rescuing all ten).

This reductivist picture can be filled out in several ways. *Invariantism* claims that all non-contrastive sentences in any context express and are reducible to propositions with the same contrast class, perhaps all available alternatives. Norcross instead endorses *contextualism*, which claims that all non-contrastive judgments in a certain context express and are reducible to contrastive propositions with *some* contrast class, but the contrast class in the reduction base varies with the context in which the non-contrastive sentence is asserted, denied, or considered. Similar kinds of invariantism and contextualism are familiar from discussions of contextualism in epistemology (Rysiew 2021; cf. 111).

Linking contexts to contrasts

Any such reductivism needs to specify which non-fundamental, non-contrastive sentences in which contexts express and are reducible or equivalent to which fundamental contrastive propositions. That relation between contexts and contrasts is supposed to be specified by Norcross’s analyses:

G-con: An action is good iff it is better than the appropriate alternative. (113)

H-con: An action harms a person P iff it results in P being worse off than s/he would have been had the appropriate alternative been performed. (117)

R-con: An action is right iff it is at least as good as the appropriate alternative. (122)

Each of these analyses has a non-contrastive left side and a contrastive right side that are supposed to be equivalent, as indicated by “iff” (which is presumably a strict rather than a material biconditional).

Several details of these analyses are questionable, but one issue is crucial here: G-con, H-con, and R-con never explicitly mention any context. Nonetheless, Norcross makes it clear (e.g. 111, 113, 117) that which alternative is appropriate is supposed to vary with the conversational context, so the notion of context is implicit in the references to the “appropriate” alternative on the right side of G-con, H-con, and R-con. The problem is that there is still no reference to context on the left side of G-con, H-con, and R-con, for their left side does not mention “the appropriate alternative.” The left side cannot be strictly equivalent or reducible to the right side if the right side refers implicitly to context but the left side does not. It cannot be true that an act is good *independently of the context* (as the left side of G-con seems to say) simply because it is better than the alternative that is appropriate *in the context* (as the right side of G-con seems to say).

This problem can be avoided by replacing G-con with this meta-linguistic claim:

G-con-meta: An utterance of a sentence of the form “The action is good” in a context expresses a proposition that is true iff the action is better than the alternative that is appropriate in that context.

Parallel modifications in H-con and R-con can be made to yield H-con-meta and R-con-meta.

Accordingly, I will assume that Norcross intended something like G-con-meta, H-con-meta, and R-con-meta. These analyses then claim strict equivalence between non-fundamental, non-contrastive sentences in everyday contexts (on the left side of the analyses) and certain fundamental contrastive propositions (on the right side of the analyses).

What changes with context?

How do these analyses handle the initial paradox? In a discussion of *Burning Building*, imagine that a context in which rescuing eight or fewer is the appropriate alternative somehow shifts so that rescuing ten becomes the appropriate alternative. According to G-con-meta, before the shift an utterance of the sentence “Rescuing nine is good” expresses the true proposition that rescuing nine is good in contrast with rescuing eight or fewer; and after the shift “Rescuing nine is not good” expresses the true proposition that rescuing nine is not good in contrast with rescuing ten. Although they appear contradictory, each sentence might seem true when one thinks of the context in which it expresses a true proposition. On the other hand, “Rescuing nine is good” will seem counterintuitive when thinking of the later context in which rescuing ten is the appropriate alternative, and “Rescuing nine is not good” will seem counterintuitive

when thinking of the earlier context in which rescuing eight or fewer is the appropriate alternative. And both sentences will seem simplistic or incomplete if one keeps both contexts in mind. This explains the appearance of paradox when one considers the sentence in the abstract without specifying any context.

This explanation might be taken to suggest that each sentence changes its truth value from before to after the context shifts. However, Norcross usually takes propositions rather than sentences to be the sole or primary truth bearers (118, 124; despite 78). If so, each sentence itself has no truth value either before or after the shift, so each sentence's (lack of) truth value does not change. But what about propositions? The sentence "Rescuing nine is good" in the context *before* the shift expresses the proposition that rescuing nine is good in contrast with rescuing eight or fewer, and the sentence "Rescuing nine is not good" in the context *after* the shift expresses the proposition that rescuing nine is not good in contrast with rescuing ten. Both of these contrastive propositions are true at both times – both before and after the shift. Of course, the sentence "Rescuing nine is not good" in the context before the shift expresses the proposition that rescuing nine is not good in contrast with rescuing eight or fewer. That contrastive proposition is false, but it stays false after as well as before the shift. Thus, neither non-contrastive sentence and neither contrastive proposition changes in truth value from before to after the shift in context. The paradox dissolves without requiring any truth values to change.

What is an alternative?

Instead of truth values, the contrast changes, but what exactly contrasts with what? Unlike other views called "contextualism" (Timmons 1998; Unger 1996), what matters to Norcross is contrasts between or among actions. When Norcross refers to "the appropriate alternative," he seems to mean an action – either an incompatible action or "the action itself" (111). Still, despite his definite article "the" in "the appropriate alternative," Norcross might not mean only one specific action. After all, rescuing nine is good or right in contrast with rescuing only eight, with rescuing only seven, with rescuing only six, and so on. Thus, someone who judges that Agent's rescue of nine is good in contrast with rescuing eight or fewer does not seem to mean only one specific incompatible action.

Norcross still might think of a non-specific action like rescuing eight or fewer as a single act that can be *the* (one and only?) appropriate alternative. Or Norcross might claim that the fundamental moral facts include only contrasts with single specific incompatible acts, so that the fundamental moral facts are (a) that rescuing nine is good or right in contrast with rescuing only eight, (b) that rescuing nine is good or right in contrast with rescuing only seven, (c) that rescuing nine is good or right in contrast with rescuing only six, and so on. Then the summary judgment that rescuing nine is good or right in contrast with rescuing eight or fewer can be reduced to the conjunction of this set of fundamental moral facts in which the alternatives are single specific acts. This summary judgment will be reducible instead of fundamental, but Norcross can admit that these contrastive judgments are not fundamental as long as the judgments that are fundamental are also contrastive.

Accordingly, I will assume that what Norcross means by "the appropriate alternative" is not "the one and only appropriate alternative" but rather "the appropriate set of alternatives." The alternatives in this set are then actions that are incompatible with each other and also incompatible with the action being assessed as good (bad, wrong, harmful) or not – perhaps plus "the action itself" (111).

This interpretation raises important questions about whether an act has to be better (or worse) than *all*, *most*, or only *some* alternatives in the appropriate set in order to be good (or bad). The answers are not obvious, because an act can be good without being the best (and bad without being the worst). Satisficers will also ask whether an act must be better than all alternatives in the appropriate set in order to be right and not wrong (21–22; cf. Snedegar 2017, chapter 4). Is an act good, bad, right, or wrong when it ties or is incomparable with one alternative in that set and is better than all others in the set? Although these problems deserve careful attention (cf. Sinnott-Armstrong 1988), here I will instead focus on another profound difficulty that arises not only for Norcross but also for other contextualists regarding not only morality but also epistemology and other topics.

Which context?

The most serious problem for Norcross is that he needs to tell us which alternatives are appropriate in a given context. Otherwise we cannot determine which sentences express true propositions. After all, according to G-con-meta, an utterance of “Agent’s rescue of nine is good” expresses a true proposition in contexts where the appropriate alternative is Agent’s rescue of eight or fewer, but not in contexts where the appropriate alternative is Agent’s rescue of ten. So, if someone utters, “Agent’s rescue of nine is good,” then we cannot know whether they express a true proposition or instead a false proposition unless we know whether the context in which they utter the sentence is one in which the appropriate alternative is Agent’s rescue of ten or instead Agent’s rescue of eight or fewer. Norcross’s contextualist analyses alone do not imply anything about which sentences in which contexts express truths or falsehoods. To specify that, he needs to add guidelines about which alternatives are appropriate in which contexts – or at least he needs to tell us how he or someone else could specify and justify such guidelines.

Norcross needs *general* guidelines, because it is not enough to announce without reason that a certain alternative is appropriate or not in a *particular* context. He makes many such pronouncements. He says about *Burning Building*, “In most conversational contexts, the appropriate alternative will be rescuing none (or perhaps one), and so the rescue of three will be judged to be good” (114). He does not give any reason for this claim, but it is not obvious. For one thing, why “most” but not all contexts? Which contexts are among the “most”? Which are in the minority? And is the appropriate alternative saving “none” or “one” – “perhaps”? I suspect that many people will think of saving four as an appropriate alternative, so they will disagree with Norcross and deny that the rescue of only three out of ten is good. Similar problems arise in other cases, including *Perot* (114; “In most conversational contexts ... the appropriate alternative will be ...”; see below), *Button Pusher* (114; “It is hard to imagine a context in which anything other than pushing ‘0’ is selected as the appropriate alternative”), and *Mary* (136; “many (probably most) conversational contexts are such that the verbal reprimand ... will be picked out as the appropriate alternative”). In all of these cases, Norcross provides little or no reason why his chosen alternative is the appropriate alternative.

The problem is that reasonable people could easily disagree about which alternatives are or are not appropriate, so Norcross needs to say not only *which* alternatives are appropriate in a context but also *why* those alternatives are appropriate in that context but not in other contexts (as well as why other alternatives are not appropriate in that

context). Imagine that one person says, “Agent’s act was good, because rescuing ten is not an appropriate alternative,” but another replies, “I disagree. Rescuing ten is an appropriate alternative, so Agent’s act was bad.” Such disputes are inevitable, and we need some way to settle them in order to determine which non-contrastive sentence expresses a truth in a given context.

In response, Norcross admits, “there may be no simple recipe” (120), and adds, “since my purpose here is simply to sketch the contextualist approach to various ethical notions, I won’t attempt a more detailed account of how context helps fix the referent of ‘appropriate alternative,’ and how salience differs from appropriateness” (122). However, we cannot assess his contextualist approach without knowing what it implies, since it is unacceptable if it implies falsehoods. We cannot know its implications without at least general guidelines regarding how contexts make alternatives appropriate or not. Nobody demands “a simple recipe” or even a “detailed account,” but we still do need to know more about what counts as a context, when two contexts are different, and which aspects of a context determine which alternatives are appropriate. We need more than Norcross supplies.

I will argue that Norcross (and other contextualists in other areas) cannot satisfy this need. Indeed, I cannot imagine how to begin moving toward “a more detailed account.” To show how difficult this journey is, I will point out several obstacles at its very beginning. Without removing these obstacles, we cannot get going on this long and winding road.

Whose context?

The first problem is that the speaker’s context, the audience’s context, and the agent’s context can come apart. After all, an audience might hear or see a recording on Wednesday about what a speaker said on Tuesday about an act on Monday. These separate contexts can differ in ways that suggest distinct alternatives as appropriate or not. Then whose context matters?

Norcross cannot think that the appropriate alternative is determined solely by the context of the agent whose act is judged. That answer would lead to invariantism instead of contextualism, since the agent has the same context regardless of who is speaking to whom about the agent’s act. Thus, Norcross must have in mind the context of the speaker, the audience, or both (though the agent’s context might be part of what matters).

Norcross suggests in some passages that the speaker’s context is what determines the appropriate alternative. For example, he refers to “the context in which the judgment is made” (111), “the situation of evaluation” (113), and “the context of utterance” (119). In other passages, however, Norcross refers to “the conversational context” (114, 117) and “the linguistic context” (115). Since the conversation involves not just the speaker but also the audience, and both use language, these references to conversational and linguistic contexts allow that what matters might be the audience’s context instead of – or in addition to – the speaker’s context.

If the audience is part of the context that determines the appropriate alternative, a problem arises. A speaker can speak to many audiences at once and often does not know who is in the audience or which “assumptions and background beliefs” (150) they share with each other or with the speaker. Consider writing a blogpost that will be read by many different groups with many different beliefs. Is the audience the union of these groups? Are their background beliefs the intersection that is shared by

each person? If so, the audience(s) will often disagree about which alternatives are appropriate, and they might not share enough to resolve their disputes.

Contextualists might try to avoid these problems by declaring that the speaker's context is the only one that matters. This move cannot be enough. After all, a single speaker can have several beliefs (as well as purposes and values) with conflicting implications about whether a certain alternative is appropriate. One speaker might, for example, believe that every agent ought to save lives when they can and also believe that we should not demand too much of people. The former belief suggests that Agent rescuing ten is an appropriate alternative, whereas the latter belief suggests that Agent rescuing ten is not an appropriate alternative. Then merely referring to the speaker's context in the abstract will not determine whether or not this alternative is appropriate.

Other problems arise when the speaker calls an act good because the speaker does not or cannot imagine any better alternative, but the agent knows more than the speaker. In particular, the agent might know that they could do something that would have been better, but they rejected that alternative for selfish reasons. Then it does not seem true that the agent's act was good, even though that judgment would be true according to G-con-meta on the assumption that the speaker's context alone determines the appropriate alternative. Indeed, the same problem persists if the audience's context matters in addition to the speaker's context. If the audience and speaker both overlook the better alternative that the agent knows, it still does not seem true that the agent's act was good.

Norcross needs to deal with these difficulties if he wants to stick with the position that the speaker's context (alone?) determines the appropriate alternative. He could give up this position, but then he needs to tell us whose context does determine the appropriate alternative and why.

Norms

Some of the preceding problems arise when someone does not consider alternatives that they ought to consider. It might seem that these problems can be solved simply by focusing on norms.

This approach cannot refer to moral norms, because moral judgments are what contextualism is supposed to be analyzing. An analysis of judgments of the general form, "The action is good," would collapse into a vicious regress or circularity if the right side of the analysis referred to the alternatives that it is good to consider. After all, considering something is an action. Still, this circularity can be avoided if the analysis refers to a different kind of norm.

One other kind of norm is statistical. Some statistical norms refer to what is *normally considered* (cf. 114). The problem is that most speakers might not normally consider alternatives that do seem appropriate. Most people do not consider giving all (or half) of their money to the needy (cf. Salwen & Salwen 2010). Nonetheless, if this oversight results from unjustified greed and selfishness, then this statistic is flimsy grounds at best for not treating that alternative as appropriate, especially when an agent bucks the trend and considers and rejects this normally overlooked alternative.

Another kind of statistical norm refers instead to what is *normally done*, but similar problems arise. Suppose that very few firefighters would rescue all ten in *Burning Building*, but only because most firefighters are unjustifiably lazy or selfish and not because rescuing all ten requires unusual stamina or courage. Then this statistic does not justify denying that rescuing all ten is the appropriate alternative. These firefighters'

failure to save all ten can legitimately be judged to be bad, even if most of them do the bad act.

Yet another kind of norm that is neither moral nor statistical is *instrumental*. We often say that someone ought to consider or do an alternative action because it will serve some purpose of theirs. But again, purposes can be bad, and then it is hard to see why they should determine whether an act is good (or bad, right, or wrong) in the sense that Norcross is trying to analyze. Moreover, a single speaker can have several purposes that conflict, especially when speaking to several audiences at once. These purposes might point toward different alternatives as appropriate, and then simply citing the speaker's purposes cannot determine which act is the appropriate alternative.

The final kind of norm worth mentioning is *institutional*. Norcross refers to norms of educational institutions when he talks about moral judgments in a philosophy class (123; cf. Lewis 1996). But these institutional norms lead to the same old problems if the institution is corrupt and immoral, like a violent gang. The gang's norms might forbid disobeying its leader, but that cannot show that disobedience is bad or that obedience is good in the sense that Norcross is trying to analyze. Of course, this problem could be avoided if institutions had to be morally acceptable, but then that restriction would reintroduce the circularity discussed above.

Norcross might reply that a different kind of norm determines which alternatives are appropriate. But which other kind of norm? Or he might reply that he does not base "the appropriate alternative" on any kind of norm. But then what makes a certain alternative appropriate or not? Or he might argue that the above circularity is not vicious. But why not? It is not clear how even to begin answering these questions.

Attention

Another way to specify which action is the appropriate alternative refers not to what should be considered or what is normally considered by most speakers but, instead, to what a particular speaker actually does consider during a particular conversation. This move is often made by contextualists in epistemology (Rysiew 2021), including David Lewis, who wrote, "No matter how far-fetched a certain possibility is, no matter how properly we might have ignored it in some other context, if in *this* context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative" (1996, 559). Following Lewis, Norcross says, "simply entertaining a counterfactual may change the context" (117) and "I mean by salience, roughly, the degree to which participants in a conversational context focus on an alternative.... Salience often plays a role in determining which alternative the context sets as the appropriate one" (120–21; cf. 136, 149).

However, the mere fact that a speaker considers or mentions an alternative cannot really be sufficient to make it the (or even an) appropriate alternative. To see why, imagine that a defense lawyer argues, "My client killed only three people when he could have killed five, so he did a good deed." This utterance mentions killing five people, and the judge and jury are legally required to consider and pay attention to what the defense lawyer says. Nonetheless, this lawyer's trick would not suddenly make killing five the appropriate alternative for judging whether the defendant's action was good or right. Norcross might reply that killing none remains in the set of appropriate alternatives even after killing five is added to that set, and an act is good only if it is better than every alternative in the appropriate set. However, almost no act that any of us ever does would be good by this high standard, since almost no actual act is better than every

appropriate alternative. Moreover, it still seems misleading at best for the defense lawyer to say, “My client did not do the best possible act, but what he did was better than many appropriate alternatives.” Why should killing five be considered an appropriate alternative just because the tricky defense lawyer mentions it?

Norcross admits problems like this when he distinguishes salience from appropriateness (121). This is why he adds, “since my purpose here is simply to sketch the contextualist approach to various ethical notions, I won’t attempt a more detailed account of how context helps fix the referent of ‘appropriate alternative,’ and how salience differs from appropriateness” (122). Despite this admission of incompleteness, as I argued, we cannot assess his theory if we do not know what it implies, and we cannot know what it implies until he tells us more about how context determines “the appropriate alternative.” It is not enough simply to refer to salience and dodge the crucial notion of appropriateness.

Antecedents

Along similar lines, Norcross sometimes suggests that what was said earlier in a conversation, either by this speaker or by someone else, can affect which action is the appropriate alternative. In one of Norcross’s examples, Perot donates up to \$1000 each month but never more, he had planned to donate \$1000 to fight drought in Somalia, he instead donates \$1000 to the homeless in Dallas, and twenty children in Somalia die who would not have died if he had donated his \$1000 to fight drought in Somalia. Norcross comments,

Suppose that, just before asking whether Perot’s donation to the Dallas homeless was good, we have been discussing his prior intention to give money to the Somalians. In this context, we are quite likely to compare the actual donation with the better alternative. On the other hand, suppose that, just before asking whether his donation was good, we have been discussing the fact that Perot has made no charitable donations at all in four of the last six months, and small ones in the other two. In this context, we will probably compare the actual donation with a worse alternative. (115)

Unfortunately, what “we will probably compare” might not be what is appropriate to compare or what affects whether it is true to call an act good. If we are talking to Perot, he cannot make it true to say that his act was good simply by saying first that he gave very little in the last few months, even if that earlier utterance makes smaller donations *seem* to be the appropriate alternative. And Perot’s political opponents cannot make it true to say that his act was bad simply by saying first that he was planning to do something even better, even if that makes the better act *seem* to be the appropriate alternative. The prior topic of conversation might affect what *seems* appropriate, but that appearance does not show that it affects either what is appropriate, which acts are good, or which moral judgments are true.

Norcross might complain that these implications of his view seem implausible to us – me and my readers (I hope) – only because the other alternatives were mentioned earlier in the current article. However, according to Norcross, what is supposed to affect which alternatives are appropriate and which sentences express truths is what is mentioned in the conversations in which those sentences are uttered. In my examples, these conversations take place between Perot, his opponents, and their audiences. We – me

and my readers – are not part of those conversations. We should be able to assess whether the utterances by Perot and his opponents in their contexts expressed a truth, even if different alternatives are appropriate in our own contexts. Otherwise, we would not be able to understand what they say. And when we do assess the truth of what they say in their context, it does not depend on what they said earlier in their conversation.

Besides, the point here is *not* about what *seems* appropriate or true to us or to anybody. Recall that the left side of G-con-meta specifies when “An action is good” expresses a truth. These principles are not about what *seems* good or what *seems* to express a truth. They are about what *is* good and what *does* express a truth. These claims, therefore, cannot be defended simply by claiming that what the speaker and audience said and heard in advance affects which alternatives *seem* appropriate to them.

Descriptions

In a related passage, Norcross claims, “A change in the description of the action might change the appropriate comparison” (114). Here appropriateness is supposed to be determined not by what was said earlier in the conversation but instead by what is said in the current judgment of the action – specifically, how it is described. In Norcross’s example just mentioned, Perot’s action could be described in two ways (among many others):

(D1) This month he *donates* \$1000 to homeless in Dallas.

(D2) This month he *switches his donation* from Somalia to the homeless in Dallas.

D1 might seem to suggest that the appropriate alternative is donating nothing, so Perot’s act was good; whereas D2 might seem to suggest that the appropriate alternative is donating \$1000 to Somalia, so Perot’s act was bad, since that alternative would have had more “comparative consequential value” (108). However, this point about descriptions runs into the same problems as before. Redescribing Perot’s act might affect which alternative *seems* appropriate and then whether his act *seems* good or bad, but Perot cannot change whether an alternative really *is* appropriate – or whether it really *is* true to say that his act is good or bad – simply by redescribing what he did. The problem again is that Norcross’s principle G-con and my revision G-con-meta are about what is good and which sentences express truths rather than about what *seems* good, true, or appropriate.

Conflicts

Norcross might reply that all of these factors – norms, attention, antecedents, descriptions, and maybe more – together determine which action is the appropriate alternative. He still needs to specify how to weigh these factors against each other when they conflict. What if the speaker and audience attend to an alternative in comparison with which their action is bad, so that alternative is appropriate according to a rule of attention; but then the speaker tries to mislead the audience by describing their action in a way that suggests a different alternative in comparison with which their action is good, so that other action becomes the appropriate alternative according to the description and the speaker’s purpose? And what if someone in the audience said something earlier

in the conversation that suggests an alternative in comparison with which an action is bad, but some relevant (statistical, instrumental, or institutional) norm suggests a different alternative in comparison with which their action is good? Then this norm conflicts with a claim that conversational antecedents determine the appropriate alternative.

Without some way to resolve such conflicts, we have no way to determine which action really is the appropriate alternative or whether it really is true to say that the act is good. This gap leaves us with no way to specify what Norcross's theory implies, so we cannot determine whether its implications are correct or even plausible.

Of course, I cannot survey all possible ways to specify how the context affects the appropriate alternative. Norcross still might have some other way to solve this problem. However, it is not enough simply to give up and say "there may be no simple recipe" (120) and "I won't attempt a more detailed account" (122), because it is not clear how even to begin developing any more detailed account. The problem is not only that we do not have simple recipe. It is that we do not know which ingredients to use or how to cook them. This problem plagues not only Norcross's contextualism but also other forms of contextualism in ethics (Unger 1996) and epistemology (Lewis 1996; Rysiew 2021; cf. 111).

Contrastivism

These difficulties are created by Norcross's claim that the context makes some alternatives appropriate and others not appropriate for assessing a non-contrastive judgment in that context. Thus, these problems can be avoided by giving up that defining claim of contextualism.

But then what is left of Norcross's views? He can keep his *contrast claim* – "Any unified theory [of good action] requires the fixing of a contrast point" (113) – if this means only that every fundamental moral judgment or fact must specify a contrast, so non-contrastive claims are never fundamental. He can also keep his *comparison claim* insofar as it ascribes reasons for alternatives on the basis of their "comparative consequential value" (108). The truth or falsity of a fundamental moral judgment that is explicitly contrastive can be determined by a comparison between the consequential values of the alternatives that are explicitly contrasted.

For example, the fundamental moral judgments in *Burning Building* might include these:

It was good and right for Agent to save nine instead of eight or fewer.

It was bad and wrong for Agent to save nine instead of ten.

These judgments are true because the consequential value of saving nine is greater than the consequential value of saving eight or fewer but less than the consequential value of saving ten. None of this mentions any context in which the judgment is asserted or believed, so it is independent of contextualism. Also, none of this calls any alternative "appropriate" or not, so it avoids the difficulties of specifying "the appropriate alternative."

Accordingly, this stripped-down theory that all fundamental moral judgments are contrastive should be understood as *contrastivism* instead of contextualism. It becomes *consequentialist contrastivism* when conjoined with the claim that the truth-value of fundamental moral judgments depends only on the comparative consequential value of the contrasted alternatives.

These claims do not yet tell us how to analyze non-contrastive moral judgments. That is where Norcross needed to refer to contexts and to “the appropriate alternative.” Consequentialist contrastivists also need to say something about non-contrastive moral judgments, because they are so common in everyday conversations.

One possibility might adapt Norcross’s analogy to indexicals (cf. Unger 1996, 163). Norcross writes, “‘appropriate’ functions as an indexical, in much the same way as ‘today’ functions as an indexical” (119). The sentence “Today is Tuesday” is true in a context in which it is asserted if and only if that context of assertion occurs on Tuesday. However, this analogy to indexicals is not quite accurate. Consider the conjunction “Today is Tuesday, and today is not Tuesday.” When the entire conjunction is asserted in a single context (not crossing midnight), both instances of “Today” must refer to the same day, which means that the conjuncts are incompatible. In contrast, Norcross needs the conjuncts in “Agent’s rescue of nine is good, and Agent’s rescue of nine is not good” to be compatible in order to dissolve the paradox in *Burning Building* in the way discussed above.

Instead of citing indexicals, contrastivists need to distinguish speaker-meaning from sentence-meaning (Grice 1968). The sentence-meaning of “The action is good” remains constant across the contexts in which that sentence is uttered (apart from certain abnormal contexts), whereas the speaker-meaning of utterances of “The action is good” varies across contexts depending on what the speaker intends to convey by that utterance in that context. The challenge is then to specify the sentence-meaning of non-contrastive evaluative sentences that remains constant across contexts and speaker’s intentions.

As with epistemic judgments (Sinnott-Armstrong 2006, 94ff.), I suggest that the sentence-meaning of “The action is good” is simply “The action is good in contrast with the relevant alternatives.” To see how this account works, recall *Burning Building*. When a speaker says, “Agent’s rescue of nine is good,” that *sentence means* “Agent’s rescue of nine is good in contrast with the relevant alternatives.” And if the speaker intends to convey that Agent’s rescue of nine is good in contrast with rescuing eight or fewer, then that is the *speaker-meaning* of that utterance in that context. Similarly, when a speaker says, “Agent’s rescue of nine is not good,” and intends to convey that Agent’s rescue of nine is not good in contrast with rescuing ten, then that is the *speaker-meaning* of that utterance in that context; but the *sentence-meaning* is still “Agent’s rescue of nine is not good in contrast with the relevant alternatives.” The speaker-meaning is what the speaker intends to convey, so it changes with speaker and context, whereas the sentence-meaning remains constant across changes of speaker-meaning.

What about the conjunction? A speaker who asserts “Agent’s rescue of nine is good, and Agent’s rescue of nine is not good” can be understood as intending Agent’s rescue of nine to contrast with rescuing eight or fewer in the first conjunct (which says it is good) and then intending Agent’s rescue of nine to contrast with rescuing ten in the second conjunct (which says it is not good). The speaker-meanings of the conjuncts are then compatible, and the speaker-meaning of the conjunction is consistent. On my proposal, the sentence-meaning of this conjunction is “Agent’s rescue of nine is good in contrast with the relevant alternatives, and Agent’s rescue of nine is not good in contrast with the relevant alternatives.” Those conjuncts might seem incompatible, but they are not. To see why, compare “This child is tall” and assume that this sentence means “This child is tall in contrast with the relevant others.” If a speaker utters, “This child is tall, and this child is not tall,” the speaker-meaning can be that this child is tall in contrast with others her age, and this child is not tall in contrast with adults,

and the sentence-meaning of the conjunction can be “This child is tall in contrast with the relevant others, and this child is not tall in contrast with the relevant others.” This conjunction is consistent because the relevant others in the first conjunct can be different from the relevant others in the second conjunct. Even if this conjunction is odd and misleading, we can understand it, and it can be true. Analogously, the conjuncts in “Agent’s rescue of nine is good in contrast with the relevant alternatives, and Agent’s rescue of nine is not good in contrast with the relevant alternatives” are compatible if the purportedly relevant alternatives in the first conjunct (rescuing eight or fewer) are distinct from the purportedly relevant alternatives in the second conjunct (rescuing ten). On this analysis, we can understand both the speaker-meaning and the sentence-meaning of the conjunction, and how the conjunction can be true, so nothing remains of the paradox.

This contrastivist account can be clarified further by distinguishing it from Norcross’s analyses – G-con, H-con, and R-con – that refer to “the appropriate alternative.” The point is not simply to replace “appropriate” with “relevant.” Indeed, Norcross himself sometimes uses “relevant” in place of “appropriate” (127). The difference is also not simply that I refer to speaker-meaning, since Norcross could say that speaker-meaning or intention determines “the appropriate alternative.” Instead, the crucial difference is that the contrastivist theory is only about what the sentences and speakers mean, so it makes no commitments about whether or when those sentences are true. If it is a fundamental contrastive moral fact that Agent’s rescue of nine is good in contrast with rescuing eight or fewer, then what the *speaker* means is true when the speaker intends to convey this fact. However, what the *sentence* means is “Agent’s rescue of nine is good in contrast with the relevant alternatives.” That is true only if some alternatives really are relevant or appropriate. That assumption is what contrastivists need to avoid in order to avoid needing to spell out which alternatives really are relevant or appropriate. All of the problems that we saw for specifying which alternative is appropriate then justify refusing to assert that any alternative really is appropriate or relevant and then also refusing to make any non-contrastive moral judgment that assumes that any alternative really is appropriate or relevant. If we cannot specify which or why alternatives are appropriate or relevant, then we should avoid asserting or implying that some are and others are not. That is why contrastivists refuse to make such claims.

Crucially, contrastivists need not assert the opposite either. They can suspend belief and refuse to assert either that some alternatives really are relevant or that no alternatives really are relevant. If so, they can avoid asserting the non-contrastive judgment that Agent’s rescue of nine is good (or bad or harmful or wrong) or not. They can also avoid denying any of these non-contrastive judgments, so they are not error theorists (109). All they assert or deny – and all they need to assert or deny – are fundamental contrastive judgments such as “Agent’s rescue of nine is good in contrast with the rescuing eight or fewer” and “Agent’s rescue of nine is not good in contrast with the rescuing ten” and similar contrastive judgments about what is bad, right, wrong, harmful, etc. This suspension of belief makes this contrastivism a kind of Pyrrhonism (Sinnott-Armstrong 2006, 103ff.).

It also clarifies the difference between this contrastivism and Norcross’s contextualism. Norcross’s analyses together with his claims that certain alternatives are appropriate imply non-contrastive facts. For example, his analysis G-con says, “An action is good iff it is better than the appropriate alternative” (113). On the next page, Norcross adds, “In most conversational contexts, the appropriate alternative [to Agent’s rescue of nine]

will be rescuing none (or perhaps one)” (114). Together these claims imply the non-contrastive moral fact that Agent’s rescue of nine is good (and, indeed, Agent’s rescue of two is good). Similarly, my suggested meta-language revisions (G-con-meta, etc.) together with Norcross’s claim that rescuing none is the appropriate alternative to Agent’s rescue of nine (114) imply that “Agent’s rescue of nine is good” expresses a true proposition in that context (and so does “Agent’s rescue of two is good”).

This commitment to some non-contrastive moral sentences expressing truths and to some actions being non-contrastively good is not just an occasional slip. Norcross seems to see it as necessary in order to avoid collapsing into an error theory (109). It is also central to contextualism, since the defining claim that “what comparisons are appropriate can change with a change in the linguistic context” (114; cf. 117, 118, etc.) implies that one comparison is appropriate in one context, whereas another comparison is appropriate in another context. Moreover, while discussing how his theory applies to the non-identity problem, Norcross commits himself to non-contrastive moral judgments, such as when he says “the risky policy is bad, because it leads to a worse world than the safe policy” (147).

These commitments reveal a basic difference between contextualists like Norcross and contrastivists. Contextualists make commitments about which alternatives are appropriate in a context and which non-contrastive sentences in which contexts express true propositions. Norcross cannot avoid these commitments simply by saying that speaker-meaning or intention determines “the appropriate alternative.” That version of contextualism still implies that certain alternatives are appropriate, but others are not, and that certain non-contrastive sentences in their contexts express true propositions, but others do not. Those implications place the burden on contextualists to specify at least general guidelines about how contexts affect which alternatives are appropriate. On the other hand, contrastivists (at least in my Pyrrhonian version) avoid all of these commitments by suspending belief about them. As a result, contrastivists do not need to provide any account of how contexts affect which alternatives are appropriate. Given the difficulties in spelling out any detailed account of appropriateness, this lesser burden is a significant advantage for contrastivism.

Advice

Suspending belief in non-contrastive moral judgments might seem to come at a cost for moral guidance and advice (cf. 148–50). Suppose a potential donor asks for advice: “Should I donate \$1000?” To this, a contrastivist replies: “I suspend belief about that, but I can tell you this: you should donate \$1000 in contrast with \$100, and you should not donate \$1000 in contrast with \$10,000.” The donor is likely to be disappointed and to view this “advice” as unhelpful at best.

Contrastivists could reply that the (quoted) contrastive judgments are still true. They are also less controversial than any simple yes-or-no answer without an explicit contrast. If one person claims and another denies that the donor should give \$1000, they could fight about it for a long time without resolution. They will be more likely to agree on the contrastive moral judgments that the donor should donate \$1000 in contrast with \$100 and should not donate \$1000 in contrast with \$10,000. Thus, contrastive judgments are useful for avoiding error and controversy.

In addition, contrastive judgments can locate the act of donating \$1000 at a particular point along the dimension from best to worst by comparing it to each alternative action along that dimension. One worthwhile goal of moral theory is to map actions

precisely in moral space relative to other actions. Telling people what to do is not the only purpose of moral theory.

Besides, contrastive moral judgments can guide people indirectly, because much of moral life concerns dimensions instead of dichotomies. The fundamental moral facts that a millionaire should donate \$1000 in contrast with \$100 and should not donate \$1000 in contrast with \$10,000 imply that, if this millionaire gives only \$1000, then observers may legitimately subject this millionaire to stronger anger and criticism (as well as weaker praise) than they inflict on otherwise similar millionaires who give \$10,000 and weaker anger and criticism (as well as stronger praise) than they inflict on otherwise similar millionaires who give only \$100. The same goes for how guilty or proud this millionaire should feel for giving \$1000. In general, we want our reactions to our own and others' actions to correlate with how bad or good those actions are. That relation can be expressed by contrastive judgments about the range of alternatives in contrast with which each act is good, bad, right, or wrong.

Thus, contrastive moral judgments can play important roles in moral theory and moral life even if they do not directly answer the potential donor's request for non-contrastive advice. Far from being useless, contrastive moral judgments can help us avoid controversies, vagueness, and disproportionate reactions in addition to paradoxes.

Conclusions

I have criticized several aspects of Norcross's contextualism and proposed a contrastivist competitor. I would be remiss, however, if I did not admit that his view and mine are very close in many ways. In particular, we both accept both the contrast claim that all fundamental moral facts are contrastive and the comparison claim that we should assess alternatives by comparing their consequential value. Where we disagree is mainly about the best way to understand non-contrastive moral judgments, which are common in everyday life even if they should play no role in fundamental moral theory. Even there we are close, though our differences are important. Norcross and other contextualists state or imply that some alternatives are appropriate and others are not, whereas I argued that we have no way to determine which alternative is appropriate or relevant, even within a particular context, so we should suspend belief about all non-contrastive moral judgments.

To be fair to Norcross, recall once again that he admits, "there may be no simple recipe" (120), and "I won't attempt a more detailed account" (122). However, I argued that this gap in his theory makes it not just incomplete but inadequate, because we cannot tell what his theory implies without the details that he fails to provide. For this reason, he needs some "more detailed account" or at least general guidelines. That is a serious problem for him because, as I argued, no "more detailed account" is available, and we do not know how to begin to develop one.

Finally, I argued that we can do moral theory and live our moral lives without non-contrastive moral judgments. This all together raises my overall question for Norcross (and for my readers): Why not give up the futile quest for "the appropriate alternative" and join me as a contrastivist?

I look forward to his reply.

Acknowledgements. I am grateful to Alastair Norcross for writing such an inspiring and insightful book; I also thank Ben Eggleston, Shelly Kagan, and Elinor Mason for joining me in this symposium; to Ben Eva, Kevin Richardson, and the Duke philosophy community for comments on an early talk about this material; and to Ben Eggleston, Shelly Kagan, and Justin Snedegar for comments on an early draft.

References

- Anscombe, G.E.M.** 1957. *Intention*. Oxford: Blackwell.
- Davidson, Donald.** 1985. Reply to Quine on Events. In Ernest LePore and Brian McLaughlin (eds.), *Actions and Events: Perspectives on the Philosophy of Donald Davidson*. Oxford: Blackwell, pp. 172–176. Reprinted in Donald Davidson, *Essays on Actions and Events*, 2nd edn. Oxford: Clarendon Press, 2001.
- Goldman, Alvin.** 1970. *A Theory of Human Action*. Princeton: Princeton University Press.
- Grice, Paul.** 1968. Utterer's Meaning, Sentence Meaning, and Word-Meaning. *Foundations of Language*, 4: 225–42.
- Kim, Jaegwon.** 1976. Events as Property Exemplifications. In M. Brand and D. Walton (eds.), *Action Theory*, pp. 159–77. Dordrecht: Reidel.
- Lewis, David.** 1996. Elusive Knowledge. *Australian Journal of Philosophy* 74, 549–67.
- Norcross, Alastair.** 2020. *Morality by Degrees: Reasons Without Demands*. Oxford: Oxford University Press.
- Portmore, Douglas W.** 2020. *The Oxford Handbook of Consequentialism*. New York: Oxford University Press.
- Rysiew, Patrick.** 2021. Epistemic Contextualism. *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2021/entries/contextualism-epistemology/>>.
- Salwen, Kevin, and Salwen, Hannah.** 2010. *The Power of Half: One Family's Decision to Stop Taking and Start Giving Back*. Boston: Mariner Books.
- Sinnott-Armstrong, Walter.** 1988. *Moral Dilemmas*. Oxford: Basil Blackwell.
- Sinnott-Armstrong, Walter.** 2006. *Moral Scepticisms*. New York: Oxford University Press.
- Sinnott-Armstrong, Walter.** 2023. Consequentialism. *The Stanford Encyclopedia of Philosophy* (Winter 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <<https://plato.stanford.edu/archives/win2023/entries/consequentialism/>>.
- Snedegar, Justin.** 2017. *Contrastive Reasons*. Oxford: Oxford University Press.
- Timmons, Mark.** 1998. *Morality Without Foundations: A Defense of Ethical Contextualism*. New York: Oxford University Press.
- Unger, Peter.** 1996. *Living High and Letting Die: Our Illusion of Innocence*. New York: Oxford University Press.