

Amorality

ABSTRACT: Actions are usually grouped into one of several moral categories. Familiar ones include the morally required, the morally permitted, and the morally prohibited. These categories have been expanded and/or refined to include the supererogatory and the “suberogatory”. Some eschew deontic categories such as the above, but nevertheless allow the existence of two comparative moral categories, i.e., the morally *better* or morally *worse* than some contextually specified baseline. At the risk of adding to the clutter, I want to explore the possibility of yet a further category, viz., amoral action, or actions that, strictly speaking, lack any moral status whatsoever. I’m interested in examining this category in part because I believe that it is plausible to say that such actions exist. In addition, I think that admitting the existence of amoral actions has additional consequences for our understanding of the nature of moral obligation, in particular the semantics of the moral ‘ought’.

Actions are usually grouped into one of several moral categories. Familiar ones include the morally required (actions that, morally speaking, one must perform), the morally permitted (actions that, morally speaking, one can, but needn’t, perform), and the morally prohibited (i.e., immoral actions, actions that are not compatible with conformity to one’s moral requirements). These categories have been expanded and/or refined to include the supererogatory (a subset of permissible actions that are morally superior to other permissible actions but that are not themselves required) and the “suberogatory” (a subset of the permissible that are morally “bad to do, but not forbidden”¹). Others have eschewed deontic categories such as the above, but nevertheless allow the existence of comparative moral categories, i.e., the morally *better* or morally *worse* than some contextually specified baseline.²

At the risk of adding to the clutter, I want to explore the possibility of yet a further category, viz., amoral action, or actions that, strictly speaking, lack any moral status whatsoever. I’m interested in examining this category in part because I believe that it is plausible to say that such actions exist. In addition, I think that admitting the possibility of the existence of amoral actions has additional consequences for our understanding of the nature of

¹Julia Driver, “The Suberogatory” in *Australasian Journal of Philosophy* 70 (1992).

²Alastair Norcross, “Reasons without Demands: Rethinking Rightness” in *Contemporary Debates in Moral Theory*, ed. Dreier (Oxford: Blackwell, 2006).

moral obligation, in particular the semantics of the moral ‘ought’.³

The plan of the paper runs as follows. In §1 I distinguish the sense of amorality with which I work from other uses of that term that sometimes crop up in moral theory. §2 offers a *prima facie* case for the existence of amoral actions. §3 poses an important problem for the possibility of amorality. Given the semantics of moral requirement and its relationship to moral prohibition, it might appear that there can be no actions that entirely lack a moral status or valence. §4 discusses a variety of potential responses to this puzzle and offers my own preferred solution. §5 responds to an objection to that solution, and §6 concludes.

1. *Amorality: A Primer*

In discussing the notion of amorality, it is important to make clear the precise concept I wish to explore. Importantly, “amoral” is used in many different ways, some of which are irrelevant for present purposes. Chief among these is the concept of an “amoral agent” or “amoralist”.⁴ Familiar from discussions of moral motivational internalism, the amoralist is a character who supposedly makes moral judgments but who feels no motivation to conform to the moral judgments she makes. I do not wish to discuss the amoralist here; the term “amoral” as I use it is intended to designate a category of actions rather than individuals.

Furthermore, “amoral” might also be used to describe the choice between two actions that is, for moral purposes, trivial or neutral. For instance, if one is faced with the choice of saving one of two strangers from immediate drowning, each of these strangers is equally deserving of being saved, they will have equally good lives after having been saved, etc., one might describe the choice between saving one versus the other as “amoral”, in other words, of no moral consequence. Though certainly related to the concept I’m interested in exploring, it is not the same issue. I’m interested in *acts*, not choices

³Throughout this paper, I will treat the notion of moral requirement and the moral ‘ought’ as interchangeable—to say that one morally ought to ϕ just is to say that one is morally required to ϕ . Some have challenged this, holding that ‘ought’ instead suggests a form of recommendation, rather than requirement. See, most importantly, Paul McNamara, “Must I Do What I Ought (or Will the Least I Can Do)?” in *Deontic Logic, Agency and Normative Systems* (Springer, 1996). However, nothing rides on this for present purposes. If one prefers McNamara’s reading of the notion of the moral ‘ought’, he or she is welcome to simply substitute “requirement”; nothing in the paper or the significance of its argument will change.

⁴Cf. David Brink, *Moral Realism and the Foundations of Ethics* (Cambridge: Cambridge University Press, 1989), ch. 3.

between acts. And in cases like these, morality may be indifferent, but each act could maintain positive or negative moral valence. Typically, saving either stranger will be morally positive—certainly morally better than not saving anyone. Thus, as I use the term, neither action is *amoral*.

Some hold that Kant identified a category of amoral actions: actions that are compatible with the fulfillment of one’s duty, but that lack positive moral worth. Performing an action that is in conformity with duty (i.e., not cheating one’s customers), but not *from* duty (i.e., not cheating one’s customers from a motive of self-interest rather than a motive of duty) is, for this reading of Kant, amoral.⁵ This is a perfectly respectable use of the term “amoral”, but it is, again, not what I mean. Even if an action is neither morally required nor immoral, it can possess a moral valence, viz., moral *permission*. Moral permission is a form of moral status: moral norms have evaluated the action in question, and have at least determined that it is not a violation of one’s duty. Amoral actions as I use the term are neither permissible nor impermissible, they neither violate one’s duty nor conform to it, they are neither morally better nor morally worse than any other action. They are simply morally unevaluated.

One further clarification. It is common to admit that moral evaluation only goes so far, at least in the following sorts of cases. First, people typically (though not always) admit that actions we *cannot* perform don’t permit of a moral valence. This seems pretty straightforward, but I go further: I claim that there are actual actions individuals can (and sometimes do) perform that are perfectly amoral, lacking in moral valence. In addition, it seems right to say that some things that people actually do are not properly morally evaluable in a sort of hum-drum way that I’m not interested in exploring here. For instance, someone with Tourette’s Syndrome might develop a facial tick; someone in a fight might suddenly flinch when noticing an incoming blow. It seems right, and boring, to say that these things don’t have any moral status. But I think it right to say that these things aren’t genuinely *actions*, at least in the morally relevant sense of the term. Morality does not evaluate facial ticks and reflexes of this kind. Thus for the purposes of this paper, I’m interested not in the existence of amoral reflexes, but rather amoral actions in the morally relevant sense of the term. Of course, there is substantial controversy surrounding what it means to be an action in this sense.⁶ I’m going to rely on the intuitive characterization

⁵Immanuel Kant, “Groundwork of the Metaphysics of Morals”, Ak. 4:397.

⁶Though this literature is voluminous, two of the standouts are Jennifer Hornsby, *Actions* (London: Routledge and Kegan Paul, 1980); Alvin Goldman, *A Theory of Human Action* (Princeton, NJ: Princeton University Press, 1970).

here without further specification; the cases I offer should fit this category straightforwardly.

2. *Are There Amoral Actions?*

At this point I've clarified the concept I wish to explore. It seems perfectly respectable *qua* concept. But why should we care about it? Why is it worth exploration? Why isn't an investigation into the concept of amorality like an investigation into chess, or some similarly cockamamie category?

Indeed, given the conceptual analysis of amorality, it may seem extraordinarily implausible to think that such actions exist. How could it be that there is some particular action open for an agent at a time, but which lacks moral valence? This action must not simply fail to be morally required or morally wrong, it must fail to be morally permissible and supererogatory, be morally *better* or morally *worse*. How could a *real action* (however one wishes to understand this notion), not simply a reflex or simple bodily movement, that we *can* perform fail to be morally evaluable or fall under one of the recognized moral categories?

However, this possibility should not be entirely mysterious. For some normative domains, evaluation is not pervasive. Take *etiquette*. Etiquette simply has nothing to say about some actions. Consider my decision while at a coffee shop to order coffee with cream. This action is neither rude nor polite; it neither conforms to protocol nor is a violation. One might say that doing so is, perhaps, more polite than some baseline. Imagine the action of flipping my barista the bird. Certainly ordering my coffee with cream is *more polite* than flashing my server an obscene gesture. But this example seems to mistake two separable actions, both of which I perform when I order my coffee with cream: the action of *ordering my coffee with cream*, and the action of *refraining from flipping the bird*.⁷ Certainly the

⁷Note: this relies on a fine-grained conception of action-individuation. See Goldman, chs. 1 and 2. Though I prefer it for ease of exposition, one needn't accept this account to make my point. Even if one accepts a coarse-grained theory, normative valence will apply to actions under particular descriptions. So even if the action of refraining from flipping the bird and ordering coffee with cream is the same, it maintains an etiquette valence under the former description, but certainly lacks one under the latter description. Etiquette, therefore, is not pervasive. I will accept a fine-grained view here, but my arguments can be translated *mutatis mutandis*.

One further point here. You might think that a "fine-grained" theory of action-individuation (or the claim that moral valence applies to actions only under descriptions) renders the argument for amoral actions too easy: if, for instance, I perform a myriad of actions at any time at which I perform *some* action (i.e., if, whenever I steal Susie's candy,

latter action is more polite than the alternative (i.e., flipping the bird). But there's nothing about the action of ordering my coffee with cream that, in itself, is more or less polite than anything else I might do. Etiquette doesn't say anything about it.

Seen in this light, however, it is much less clear that amoral actions cannot exist. Everything hinges on whether or not morality's evaluation includes all potential actions in a way that etiquette's does not. Some believe it does. Scheffler, for instance, writes: "all human conduct is in principle morally assessable."⁸ If this is correct, there is no category of amorality. But I think there are reasons to second-guess Scheffler's confidence. Two arguments follow, focusing on roughly the same problem.

2.1. *Moral Requirability*⁹

I also perform the separate action of moving my arms in a particular way, of omitting to buy milk at the store, etc.), surely *some of them* (i.e., moving my limbs in such-and-such a way) will not be morally evaluable. And hence, or so it would seem, it is utterly trivial that there should be amoral actions. However, even if my argument here is committed to a fine-grained theory of action-individuation, this does not entail that the existence of amoral actions is uninteresting or obvious. Take, for instance, my stealing of Susie's candy. Surely that action is morally evaluable. But it seems unwarranted to me to say that moving my limbs in such-and-such a way is not morally evaluable without *much* more information. Indeed, depending on the circumstances, it seems perfectly acceptable to describe that action—considered in and of itself—as *permissible*. However, if circumstances are otherwise, it could also have other moral valences: moving my limbs in such-and-such a way, after all, is the sort of thing that could generate good or bad consequences (especially in this case), to name but one significant moral factor. It thus seems to me not so implausible to believe that doing so could even be *required*. Furthermore, it certainly does not lack any of the relevant disqualifying factors (discussed below) that paradigmatically un-requirable actions lack. Of course, this is just one example, but the general point holds: the existence of amoral actions is not simply guaranteed or a trivial matter given my assumptions here. However, even if I'm wrong about this, even if the existence of amoral actions *is* trivial, their existence has important upshots when it comes to the semantics of moral requirement, upshots I explore beginning in §3.

⁸Samuel Scheffler, *Human Morality* (Oxford: Oxford University Press, 1992), 25.

⁹Classically, many have thought that the concept of an *excuse* removes the action so excused from moral consideration. Excuses, it is said, are different than justifications: while justifications serve to show that the action in question, though it may have had some bad qualities, is nevertheless justified, an *excuse* does not serve to justify the action in question. (See, classically, J. L. Austin, "A Plea for Excuses" *Proceedings of the Aristotelian Society* 57 (1956-7), 2-3.) As Austin writes, in offering justifications for an action, "we accept responsibility but deny that it was bad"; in offering excuses, "we admit that it was bad but don't accept full, or even any, responsibility," (Austin, 2).

However, though I'm compelled to believe that *some* actions for which excuses are available are those that lack moral valence, further argument is required to show that actions for which there is a legitimate excuse possess no moral valence at all. The most

Consider the “problem of moral requirability”. The problem runs as follows. In any set of circumstances, a particular moral agent will face a number of action options, actions that the agent can perform. The problem of moral requirability, then, is the problem of determining which actions that one can perform are eligible to become the object of a moral requirement for this agent: it is the problem of determining, for this person, which actions are *morally requirable*.

One obvious response to the problem of moral requirability is to identify the set of morally requirable actions just with those actions one can perform. In other words, in considering which actions are eligible to be required, we look at those one can perform, and we assign them relative deontic valences (including requirement) based on their relative moral quality. However, there are well-rehearsed reasons to believe that the problem of moral requirability should not be solved in this way.¹⁰ For instance, consider the fact that you and I *can*, in the strictest sense of “can”, cure cancer.¹¹ All we would need to do is to sit down at our keyboards, type the right characters in the right order, and mail the result to a top medical journal or some recognized medical authority.¹² We *can* do this. Of course, we *won't* do this, for many reasons. We won't do this, most obviously, because we don't

important problem is that, as Austin notes, an excuse is not an account of the moral status of an action, but rather an account of the relationship between the agent and the action who performed it. In “refusing to accept responsibility”, one, in effect, denies that one could be sensibly *blamed* for a particular action. This seems right. But it remains possible to say that in the case of excuse, the action one performs is still *wrong* or *immoral*—it is simply the case that this morally wrong action is not *blameworthy*. The connection between moral wrongness and blameworthiness is fraught with controversy, of course. But we typically believe that wrong action is not blameworthy if is performed under conditions for which excuses are typically appropriate: ignorance, or lack of a free or autonomous will. (See Stephen Darwall, *The Second-Person Standpoint* (Cambridge, MA: Harvard University Press, 2007), 92-4.) Thanks to [...] for urging my attention to this issue.

¹⁰This problem is discussed in much more detail in [*Author's Forthcoming Paper*].

¹¹One might legitimately ask to what sense of ‘can’ I refer here. As I understand it, it is the sense of ‘can’ that is implied by ‘did’. Alfred R. Mele identifies this as a “simple ability” (see “Agent’s Abilities” in *Noûs* 37 (2003), 448). If I ϕ -ed at t , it follows that, at t , I can (I have the “simple ability to”) ϕ . Of course, this sense of ‘can’ holds the laws of nature and my own physical and mental capacities fixed. I cannot lift a two-ton boulder with my bare hands. But I can, among other things, pick the winning lottery numbers, beat Gary Kasparov, write the Great American Novel, etc. None of these are beyond my physical or mental capacities. They merely require *luck*.

¹²Cf. Eric Wiland, “Monkeys, Typewriters, and Objective Consequentialism” in *Ratio* 18 (2005).

know what the right characters are. But it would seem sensible to believe that *if* one of the actions eligible to form the object of a moral requirement for me includes *curing cancer*, a moral theory would have to be crazy not to select that action as the one that, in fact, *is* required. But of course this is absurd. There are plenty of actions that we can perform but that we almost certainly won't perform for reasons that render a moral requirement of that action implausible. For instance, it could be that opening a safe to which one does not know the combination produces an astoundingly good set of results. Or it could be that one has the power to rescue one's friends from calamity if only one can beat their captor—who happens to be a Grand Master—at chess. Surely one *can* beat the Grand Master; all one has to do is make the right moves in the right order. But while one certainly can beat the Grand Master, it's just overwhelmingly unlikely that one *won't*.¹³

Virtually everyone agrees, then, that actions that one can perform *merely* in this maximally expansive sense of 'can' cannot be eligible for moral requirement. So much, it seems to me, is simply platitudinous. What is required, however, is an account of what is missing in such actions that *renders* them ineligible for moral requirement. Different theorists have diagnosed this problem in different ways, and have searched for alternative accounts of actions that are eligible to be morally required or the object of a moral 'ought'. Ecumenically, we can at least say that any such action will maintain a set of "disqualifying factors". These factors will typically include relations borne between a particular agent and the action in question, and might include, for instance, the fact that the relevant agent has positively no clue *how* to perform the action, or is unable to perform the action as a result of reasoned deliberation, etc.¹⁴ The project of identifying precisely the disqualifying factors is, of course, controversial. But whatever these factors are, they must explain the relevant datum, viz., that most people cannot be morally required to cure cancer, despite the fact that they *can* do so. But however this dispute is resolved, the central point (for my purposes) remains. For virtually any person at virtually any time, there are genuine action options (actions that he or she can perform) that are not eligible to be morally required of that person at that time.

This problem also extends to the category of immorality. There will be

¹³For similar cases, see Frances Howard-Snyder, "The Rejection of Objective Consequentialism" in *Utilitas* 9 (1997); Dale Miller, "Actual Consequence Act-Utilitarianism and the Best Possible Humans" in *Ratio* 16 (2003).

¹⁴I have argued elsewhere that these disqualifying factors generally center around the fact that an action cannot be performed as a result of deliberative agency, but I'll leave this aside here. See [*Author's Forthcoming Paper*], op. cit.

some actions that an agent can perform, but which seem ineligible for the category “morally prohibited”, in virtue of the relationship to which an agent stands to the potential *not* to perform the action. Imagine, for instance, that Phil stands before a safe, and that opening the safe will prevent a very serious calamity. Phil, however, does not know the combination of the safe, and hence does not know how to open it. It would seem inappropriate to say that *opening* the safe is morally required, given that he doesn’t know the combination (though, strictly speaking, he *can* open the safe, by entering the right numbers in the right order). But it would also seem inappropriate to say that *failing* to open the safe is *immoral*. Given the presence of relevant disqualifying factors, the actions of opening the safe or not are not eligible to be morally required or immoral.

I think the above set of claims are reflective of moral experience; some actions are just ineligible to be morally required despite the fact that one can (if one is, e.g., really lucky) do so, and hence refraining from such actions is ineligible to be morally prohibited. But this isn’t yet enough to describe such actions as *amoral*. Note that failing to be *required* does not entail failing to be *evaluated*. A number of potential evaluations could remain; these actions could be morally permissible, as well as morally better or worse than potential alternatives.

However, *if* we allow a potential action ϕ cannot be morally required on the basis of various disqualifying factors, this entails that ϕ -ing must be amoral. The key is the relationship between moral requirement and moral *reasons*. As I understand the latter concept, a moral reason is the evaluative currency of morality: it is a fact that helps to determine whether a particular action is morally better, morally worse, or maintains some other moral deontic status, etc. So to allow that ϕ -ing is morally evaluable is to allow that moral reasons apply to it, and will count in favor of, or against, various deontic valences possessed by the action. But if we allow that moral reasons will apply to opening the safe, it seems hard to say how doing so could fail to be *morally required*. Any plausible theory, other things being equal, will require Phil to open the safe if doing so prevents a very serious calamity at basically no cost—on the assumption that moral reasons apply to so doing. Thus, if moral reasons tell in favor of opening the safe to any degree (which they must if this action is to be evaluated at all), it is hard to see how these moral reasons could *fail* to generate a moral requirement. Hence to claim that Phil’s opening the safe is not morally required—which, it seems to me, it surely isn’t—we must say that reasons don’t apply to it, or that any potential reasons are silent given its failure to be morally requirable. Parallel reasoning applies to Phil’s potential failure to open the

safe. Given that to do so is to fail to prevent a serious calamity at basically no cost, if it is evaluable, it is surely immoral. But it isn't immoral. And so it cannot be evaluable.

This shows that, e.g., Phil's action of opening the safe cannot be morally evaluated when, were it to be evaluated, it would pass whatever threshold of comparative moral quality to make it the required alternative. And hence, in cases in which two facts hold, amoral actions exist: first, in which a particular action ϕ , for agent A, and time t displays the relevant disqualifying conditions and, second, in which ϕ -ing at t , were it to be evaluated, would be required of A. This is enough, I think, to establish that amoral actions exist.

However, I don't want to leave the matter there. The field of amoral actions extends far further. Indeed, I argue, *any* action that displays the relevant disqualifying conditions in whatever context must be amoral. To see this, note two things. First, whether a particular action *would* be morally required were it to be evaluated depends on the alternatives to that action available in the circumstances. This seems straightforward enough: whether Brutus' killing of Caesar was morally required or not depends on the alternatives available to Brutus and the moral quality of those alternatives. If killing Caesar was the morally best among those alternatives, surely killing Caesar is required of Brutus.¹⁵ Second, whether or not a particular action is subject to moral evaluation in a given case is *not* dependent upon the *alternatives* to that action or their evaluability.¹⁶ Whether, in circumstances c , A's curing of cancer is morally evaluable cannot depend on what other actions are available to A in c . One might put this (seemingly straightforward) thought in the following way:

¹⁵Notice that this is not incompatible with a so-called "satisficing" approach to moral rationality. In any determination of whether a particular action is morally good *enough* in a particular case (which is sufficient, on a satisficing approach, to render it morally acceptable), this fact will surely depend on what alternatives are available. If, for instance, I have only the option to destroy Great Bend, Kansas or Paris, France, the morally "good enough" option is surely to destroy Great Bend. But if my options include destroying no city at all, destroying Great Bend is surely *not* morally good enough.

¹⁶Of course, one is quick to point out that the moral *valence* of ϕ -ing might change depending on the alternative actions available. If ϕ -ing causes a low-grade amount of harm, but the only alternative is to cause a medium-grade amount of harm, ϕ -ing would be morally required. But this would change if an alternative includes causing no grade of harm whatever. But this is perfectly sensible: in this case, the set of actions one has open changes, and so the rank-ordering of those actions will change. But it is difficult to see how a particular action could or couldn't be morally *evaluable* depending on whether some other action is or isn't available, and whether it is or isn't morally evaluable.

Independence of Evaluable Alternatives (IEA): whether ϕ is morally evaluable for a person A in circumstances c is independent of the availability and moral valence of alternative non- ϕ actions in c .¹⁷

But if we accept IEA, it would appear that any action that maintains the relevant disqualifying factors cannot be evaluated under *any* circumstances—*whether or not, were it to be evaluated, it would be required*. This is because (a) it cannot be evaluable if *to* evaluate it would be to require it *and* (b) whether or not to evaluate it is to require it depends on the alternative actions available in those circumstances. Thus to say that the relevant action can be evaluated so long as *to* evaluate it would be to not render it morally required violates IEA: it holds that an action can be morally evaluated or not depending on the alternative actions available in a particular set of circumstances. Hence, or so it would seem, the only way to guarantee that actions maintaining the relevant disqualifying factors will never be morally required is to outlaw moral evaluation of such actions in any circumstances. Moral evaluation is, therefore, *not* pervasive: it does not evaluate actions that, for particular agents, maintain the relevant disqualifying factors.

2.2. Regret¹⁸

The problem of moral requirability succeeds in defending the existence of amoral actions for views that accept the existence of moral requirements. But some views deny this. For instance, according to a *scalar* view, there are no moral requirements: actions are simply rank-ordered from best to worst.¹⁹ Because this view doesn't admit of requirements, it need not assign a valence of amoral to avoid having actions that display the relevant

¹⁷One question is worth discussing here. IEA can also be stated as follows: take two sets of circumstances, c and d . According to IEA, any action in c that is evaluable in c must also be evaluable in d if c and d differ only in the potential available actions. But, or so it would seem, this requires me to say that a particular action ϕ in c is *identical* to another action ψ , in d . But *surely* these actions are not *identical*. After all, they bear different relations (including being present in different circumstances). However, IEA can be easily restated as follows: according to IEA, the evaluability of any iteration of a particular *act-type* cannot vary *simply on the basis of* the alternatives available to this iteration. In other words, any iteration of the act-type “killing Caesar” cannot be evaluable simply as a result of the alternative actions available in the set of circumstances in which that iteration is available (though it can be evaluable or not depending on the presence or absence of, e.g., relevant disqualifying factors in such circumstances). (Goldman, 12.) Thanks to [...] for urging me to discuss this issue.

¹⁸Thanks to an anonymous reviewer for prompting me to include this argument.

¹⁹Norcross, op. cit.

disqualifying factors classified as morally required.

This proposal (and others like it²⁰) would still allow the open alternative of curing cancer to be “morally best”, “morally optimal”, or “supported by strongest moral reasons”, even if not strictly speaking required. But while this proposal is not *as* implausible as the claim that these actions are morally required, it is bad enough.

Here’s an argument for this claim. Consider the following principle concerning the appropriateness of a particular attitude—*regret*—in response to not having performed specific actions. It runs like this: if an action is morally best (or is supported by strongest moral reasons, or is morally optimal, etc.), it is, other things being equal, appropriate to feel regret for not having performed the action in question. This is a *very modest principle*. By “appropriate” here, I do not mean that the morally upright agent will *always* feel regret for not having performed the morally best action, or that someone is somehow a morally imperfect agent for not feeling such regret. I mean only that regret would not be considered “out of place”. I also do not mean to adopt a position that accepts an *analytic* link between the appropriateness of such an attitude and failing to perform the morally optimal action. Nor do I mean to say that it is appropriate to *blame* those who fail to perform the morally optimal action or the action supported by strongest moral reason (indeed, such a claim would be manifestly implausible especially if we accept the existence of the supererogatory). I simply mean that whenever an action is morally best, morally optimal, or supported by the strongest balance of moral reasons, it would not be out of place to regret failing to perform it even if it’s not *actually* the subject of a moral obligation.

One might argue that the principle I propose is too strong. Instead, one might hold that regret is appropriate only if one is also *required* to perform the action for which one has strongest moral reason. (In other words, moral regret is appropriate only if one has behaved immorally.) But, first, this is implausible especially for a scalar view. If it is appropriate to regret failing to perform an action only if that action is required, and if no actions are ever required, it is *never* appropriate to regret failing to perform an action, no matter how morally favored. This is an absurd result. But, second, this claim is implausible on any view. Take, for instance, failure to perform a *supererogatory* action, e.g., my refusal to go out of my way to

²⁰This includes a potential view according to which moral requirements exist, but according to which the relevant disqualifying factors block moral reasons (which remain applicable to such actions) from *generating* a moral requirement. Thanks to an anonymous reviewer for this proposal.

rescue a stranger's hat that had blown away in the wind.²¹ Is it plausible to say that my regret for not having performed this action is *out of place*? Inappropriate? Surely not! I could certainly say, with perfect uprightness, that I regret not doing it, insofar as it would have been a perfectly nice thing to do, the morally best thing to do, even if I wasn't strictly morally required to do it.

The appropriateness of regret for failing to have performed the morally best action is, as I interpret it here, undeniably plausible. However, it is surely the case that much of what drives our revulsion to describe, say, the action of curing cancer as not being morally requirable is the thought that it would be utterly inappropriate to feel regret for not having done so! If I regretted, say, not curing cancer all the time, I could certainly, again with perfect uprightness, be told to "snap out of it"! (Actually, strictly speaking, I don't even need a principle as strong as the one mentioned above. I could say, instead, that it is not out of place to regret not performing an action, when that action is not morally required, if that action is overwhelmingly morally better than any other possibilities one could have performed, and it was possible to perform it at basically no cost. This principle is undeniable, and would just as easily deliver the needed result, at least in the case of curing cancer.) But if this is correct, there is further reason to believe that actions that maintain the relevant disqualifying factors should not be morally evaluated. To evaluate them leaves open, depending on the alternatives, the possibility that such actions would or could be morally best, and hence appropriate to regret having failed to perform. And hence they should not be evaluated not just in contexts in which to evaluate them would be to render them morally best but also, given IEA, in any other circumstances, as well. (For brevity, I'll focus on views that accept the possibility of moral requirements for the remainder, although some remaining objections, along with my replies, can be put in terms of moral optimality with no loss of effect.)

2.3. Objection: Can't Curing Cancer be Immoral?

One might object to the reasoning on display here. Sometimes it may seem right to say that curing cancer even with these disqualifying factors could be *immoral* (and hence evaluable). For instance, imagine that one has the

²¹I owe this example to James Dreier, "Why Ethical Satisficing Makes Sense and Rational Satisficing Doesn't" in *Satisficing and Maximizing: Moral Theorists on Practical Reason*, ed. Byron (Cambridge, 2004), 147-8.

option to push a button to ensure high-quality lives for all humanity until the end of time. Instead, one simply sits down at one's computer, punches a bunch of random keys, and by sheer lucky guess, cures cancer. Wouldn't we say that curing cancer, in this case, is *immoral* (or at least morally worse, or morally negative), given the alternative? In response, I think the answer is no. Again, I think this case mixes up two separable actions, one evaluable, the other not. It is plausible to say that we could morally criticize someone for *omitting to ensure high quality lives for all when they had the chance*. This action does not maintain the relevant disqualifying factors. But this is different than saying that the action of *curing cancer*, which is distinct, is itself evaluable.²²

And given the reasoning on display above, it seems right to say that it cannot be evaluable, given that it cannot be requirable. This is not to say that there is nothing morally criticizable in what someone does in that case. One natural reaction to this case might be that despite the fact that one didn't press the welfare button, which was wrong, *at least* typing random keys wasn't a waste: one *cured cancer* by doing so; but this just illustrates that we will generally not morally evaluate the *curing of cancer itself*, but rather will treat the curing of cancer as a morally salient *effect* or *result* of an action that is, genuinely, evaluable, viz., punching random keys rather than pressing the welfare button. *Curing cancer* will not be the focus of moral criticism.

A similar objection arises in this context.²³ Consider the following plausible principle of reasons:

Incompatible: if r is a reason for x to ϕ at t , and if x 's ψ -ing at t is incompatible with x 's ϕ -ing at t , r is a reason against x 's ψ -ing at t .

Incompatible seems plausible: if I can't have a hamburger *and* a hot dog, any reason in favor of eating the hot dog would also seem to be a reason against having a hamburger, and so forth. With this in mind, consider the argument of the last paragraph. *Ex hypothesi*, pressing the high-quality-lives button is incompatible with curing cancer. But if there are a number of moral reasons in favor of pressing the button, as would surely seem to be the case, then *Incompatible* implies that these self-same reasons are also reasons *against* curing cancer, and hence it would appear, at least in some circumstances, that curing cancer can be evaluable (and potentially immoral).

²²Cf. Goldman, ch. 2.

²³Thanks to an anonymous reviewer.

Despite its plausibility, I resist *Incompatible*. Indeed, there is independent reason to reject it. This is because it doesn't restrict the actions to which moral reasons could apply *at all*. Take, for instance, the action of my reversing the direction of time for the purposes of stopping the holocaust. Doing so, let's say, is incompatible with my eating a hamburger in at least this straightforward sense: were I to reverse time to stop the holocaust, I would not eat my hamburger. But this does not entail that any reason in favor of eating a hamburger is also a reason *against* going back in time to stop the holocaust. This is because, plausibly, we *don't evaluate actions one cannot perform*. To accept *Incompatible* one would have to revise it to avoid this result. Instead, one would have to accept something along the lines of:

*Incompatible**: if r is a reason for x to ϕ at t , and if x 's ψ -ing at t is incompatible with x 's ϕ -ing at t , and if x can ψ at t , r is a reason against x 's ψ -ing at t .

This would avoid the problem just cited. But notice that to restrict the application of *Incompatible** to actions that one *can* perform just is a choice on our part: it is a choice made given the plausibility or implausibility of applying moral status to actions one cannot perform. Notice that one needn't accept my intuitive judgment that all acts one cannot perform are not evaluable; the general point stands if one more restricts *Incompatible* in a different way. And the general point is: *Incompatible* must be restricted, and it is restricted simply on the basis of our first-order judgments concerning the class of actions to which moral reasons (or reasons of any other domain, for that matter) legitimately apply. But I have already given strong reason to believe that applying moral status to actions that possess the relevant disqualifying factors is itself implausible. And hence, if that argument succeeds, there is little reason to accept *Incompatible** (or some other restriction) and not

*Incompatible***: if r is a reason for x to ϕ at t , and if x 's ψ -ing at t is incompatible with x 's ϕ -ing at t , and if x can ψ at t , and if ϕ does not maintain any of the relevant disqualifying factors for x at t , r is a reason against x 's ψ -ing at t .

But obviously *Incompatible*** causes no problem at all for the existence of amoral actions, at least as I defend them here.

2.4. *Objection: Objective/Subjective*

For an appeal to the problem of moral requirability to be useful in defending the existence of amoral actions, one must share the judgment that

actions that possess the relevant disqualifying factors cannot be morally required (or morally optimal—I’ll drop this for brevity here). But perhaps there is reason to rethink that judgment. After all, many moral theorists (including myself) are tempted to accept a distinction between *objective* and *subjective* moral requirements. While the precise account of this distinction will differ,²⁴ it is generally accepted that subjective moral requirements will be relativized to a given agent’s beliefs—one subjectively ought to perform that action that, were one’s beliefs correct, one objectively ought to perform. But we may be tempted to accept that one can be *objectively* morally required to cure cancer so long as we are not stuck with the conclusion that one is also *subjectively* morally required to cure cancer—especially so given that we typically assign *blame* for immoral actions not on the basis of its objective immorality, but on the basis of its immorality given the agent’s (reasonable) beliefs. And given that one doesn’t know how to cure cancer, it may seem plausible to say that one is not *subjectively* required to cure cancer; merely objectively so.²⁵

This proposal is in error.²⁶ First, it remains implausible to say that we are objectively required to cure cancer, even if there exists a *subjective* obligation that we can fulfill by not doing so. Indeed, if the objective ‘ought’ tells us that we ought to cure cancer, it is a little difficult to see what role, if any, the objective ‘ought’ should play in our lives. But second, and more importantly, there is no reason whatsoever to believe that we are not *subjectively* required to cure cancer, assuming that we can do so, and assuming that subjectively morally requirable actions include any that we *can* perform. After all, it is certainly no mystery that the moral value (or, perhaps, subjectively expected moral value) of curing cancer is extraordinarily high; certainly higher than the moral value (or subjectively expected moral value) of doing anything else. We all believe this, and we believe this with ample justification; furthermore we all believe, or should believe, that we *can* cure

²⁴For a number of different proposals, see David Brink, “Prudence and Authenticity: Intrapersonal Conflicts of Value” in *Philosophical Review* 112 (2003), 220; [*Author’s Published Paper*]; Allen Gibbard, “Truth and Correct Belief” in *Philosophical Issues* 15 (2005); Niko Kolodny and John MacFarlane, “Ifs and Oughts” in *Journal of Philosophy* 107 (2010), 117; Michael Zimmerman, *Living With Uncertainty* (Cambridge: Cambridge University Press, 2008), 5; Derek Parfit, *On What Matters, Volume One* (Oxford: Oxford University Press, 2011), 163.

²⁵Thanks to [...] for suggesting this response. Indeed, this response reflects the dialectic of the problem of moral requirability. This problem is usually stated as an objection to *objective* versions of act-consequentialism. Cf. Wiland, op. cit. Howard-Snyder, op. cit.; Miller, op. cit.

²⁶I have discussed this matter at some length in [*Author’s Forthcoming Paper*.]

cancer. On no plausible account of the nature of the subjective ‘ought’ am I not required to cure cancer, given that I believe, rightly (or so we are assuming for the sake of argument) that its moral value trumps anything else I might do, and that I believe (rightly) that I can do so.

I suspect that the suggestion that I am not subjectively required to cure cancer rests on the error of confusing the moral value of an action with the moral value of the means to that action. To say that I am not subjectively required to cure cancer is to treat the subjective moral value of curing cancer as deriving from the subjective moral value of *sitting down at a computer and typing random characters*, which is the only means by which I could cure cancer (and which, obviously, has a pretty low subjective moral value). But notice that these are *distinct acts*. Even if typing random characters does not have a particularly high subjective moral value, curing cancer sure does, and hence we can still be required to cure cancer (by *whatever* means).²⁷ Thus there is reason to believe that we should exclude some actions from the category of subjective moral requirability. Hence, given that non-requirability entails amorality, there are subjectively amoral actions. But it is hard to see why we should accept subjectively amoral actions without also accepting objectively amoral actions—especially given that it remains implausible to hold that we objectively ought to perform actions that we can perform only as a matter of sheer luck. And hence there are reasons to accept amoral actions at both the subjective and objective level. Given this, I hereby abstract from this distinction for the remainder.

2.5. *Objection: Justifying Reasons*

Here’s a restatement of the first argument for amoral actions. Moral evaluation of any action in any context depends on there being reasons in favor or against performing the relevant actions. But if moral reasons apply to a particular action, it would appear that, in at least some contexts, these actions (which nevertheless maintain the relevant disqualifying factors) can themselves be morally required (or immoral). One could say that they are not morally evaluable only in cases in which were they to be evaluated, they would be morally required. But this response has two problems: it continues to allow the existence of amoral actions, and (more importantly) violates IEA. And hence actions that maintain the relevant disqualifying

²⁷Note that even if we adopt a coarse-grained theory of action-individuation, we cannot escape subjective moral obligations to cure cancer (and to act in other disqualified ways): given its extremely high subjective moral value, I am surely morally required to perform the single act of typing keys in the order necessary to cure cancer.

factors cannot be morally evaluated.

As should be clear from the above argument, whether actions that maintain the relevant disqualifying factors, and are hence not *requirable*, can be morally *evaluable* turns on the nature and function of moral *reasons*. So far I've been writing as if all moral reasons have the power to require actions. In other words, once moral reasons in favor of a certain action ϕ are strong enough, or strong enough in comparison to the moral reasons in favor of an alternative action ψ , ϕ -ing is required. But, as some have argued,²⁸ moral reasons can have different functions. It could be, for instance, that some moral reasons have, at most, the power to justify actions without also maintaining the power to *require* actions. And if this is correct, we might say that reasons that count in favor of actions that maintain the relevant disqualifying factors maintain only justifying strength. No matter *how strong* the reasons in favor of curing cancer are, on this proposal, they could never require doing so, given that they fail to have the power required to do so.

However, I think this suggestion is unattractive. Note that justifying reasons still maintain at least some, indirect, relationship to moral requirement. If an action is supported by justifying reasons, but justifying reasons that *are not sufficient to justify* (given the circumstances in which it is placed), this action is thereby morally *unjustified*, and hence one is morally required *not* to perform the action. In addition, if a single action passes the relevant threshold of sufficient justification, this action necessarily will be required: it is the only morally justified action. But if we allow that justifying reasons count in favor of actions that maintain the relevant disqualifying factors, it becomes inevitable that *some* actions that maintain the relevant factors will not be favored by justifying reasons to a sufficient degree, given potential alternatives, and hence will be morally *unjustified*. But this itself is problematic. Take again Phil's case. It certainly seems wrong to say that his failure to open the safe is immoral. But, if evaluable, it seems wrong to say that failing to open the safe could possibly be justified by *sufficient* justifying reason. (After all, this failure results in some horrible calamity.) But if this is right, given that actions for which one lacks sufficient justifying reason are not justified, and hence immoral, one must say that Phil's failure to open the safe, if evaluable, is immoral.²⁹ But this is no better than saying

²⁸Most notably, Joshua Gert in *Brute Rationality* (Cambridge: Cambridge University Press, 2004), ch. 2. See also Douglas Portmore, "Are Moral Reasons Morally Overriding?" in *Ethical Theory and Moral Practice* 11 (2007).

²⁹Even if we reject deontic valences for actions that possess disqualifying factors, a roughly similar problem remains; if this action is not supported by strong enough justifying reasons, even if not immoral, it would certainly appear eligible for appropriate regret;

that, for Phil, opening the safe is morally required.

I think a summary point should be made here. It may be possible to further chip away at the set of amoral actions, given certain maneuvers with respect to when particular actions could be morally evaluated based on a number of different extrinsic factors. But I think two things have been shown. First, it seems right to say that even if such chipping away is possible, amoral actions exist. But, second, it also seems correct to say that the *best explanation* of the various cases I've brought up, and the considered judgments I, and I hope the reader, maintain, is that actions that maintain the relevant disqualifying factors (such as being an action someone fails to know how to perform, etc.) are, broadly speaking, not susceptible to moral evaluation.

3. A Puzzle: Moral Obligation and Moral Valence

So far I have offered two reasons to believe that moral evaluation is not pervasive. I think there are some actions that just don't permit of moral evaluation: at the very least, actions are amoral when they are among those actions an agent can perform, but that maintain the relevant disqualifying factors.

But even if we want to accommodate the existence of amoral actions, it's not clear that we can. The problem appears to be the semantics of moral obligation. Imagine at a particular time t , there are three actions one can perform: $\{\phi, \psi, \text{ and } \pi\}$. ϕ maintains the relevant disqualifying factors. ψ and π are both morally evaluable: ψ -ing is clearly morally superior to π . And hence it seems right to say that ψ -ing is morally required. (Let's say, for instance, that ϕ -ing is the action of curing cancer, ψ -ing is the action of costlessly promoting some important moral good, π -ing is the action of refraining from doing so.) But imagine that rather than ψ -ing, one ϕ s at t , and in so doing performs an *amoral* action. But *ex hypothesi* one is morally required to ψ , and hence in ϕ -ing, i.e., *not ψ -ing*, one has behaved contrary to one's moral obligations. Given that it is immoral to fail to conform to one's moral obligations, it would appear that, in ϕ -ing, one has behaved immorally. But *immorality* is not *amorality*. And hence ϕ -ing, at t , is not amoral.

This puzzle can be put in more general terms. For the sake of terminological expediency, refer to the set of actions someone can (has the "simple ability" to) perform as S . Refer to the set of actions that someone can (has

Argument Two therefore applies straightforwardly.

the “simple ability” to) perform, but that also are morally evaluable, as R . Now assume that at time t , agent x faces the possibility of performing an amoral action. At this time, then, there will be a distinction between the members of S and the members of R for x . But, for the members of R , there will be a subset—call this R_p —of R the members of which pass the threshold of moral permissibility. And hence, given that one performs a morally permissible act if and only if one conforms to an action one is morally required to perform, one will be required to perform some member of R_p . But given that one behaves immorally if and only if one fails to conform to one’s moral obligations, if one fails to perform a member of R_p , one will behave immorally. But any non- R member of S will, obviously, not be a member of R_p . And hence, in performing a non- R member of S , one behaves in an immoral way, and hence not an amoral way.

Schematically, this puzzle runs like this:

1. Assume that, in any choice circumstance, a proper subset of R (R_p) is constituted by those members of R that are morally permissible.
2. One is morally required to perform a morally permissible action.
3. Because amoral actions are, *ex hypothesi*, not members of R , no amoral action is a member of R_p .
4. Hence to perform an amoral action ϕ is to fail to perform a morally permissible action.
5. Hence to perform an amoral action ϕ is to fail to conform to a moral requirement.
6. Hence, ϕ -ing is immoral.
7. Hence, ϕ -ing is not amoral.

4. Responses

To defend the existence of amoral actions, it must be the case that the argument framed above goes wrong somewhere. I’ll explore three potential responses in this section.

4.1. Amorality as Conformity to Moral Obligation

First, one might deny (4). One might claim that simply because one performs an action that is not morally evaluable (i.e., not a member of R), this does not entail that one does not *also* perform a member of R_p . Take, for instance, the act of not physically assaulting the person sitting next to me.

This action is surely morally required. But imagine that instead of physically assaulting the person next to me I whip out my laptop, type a bunch of random keys, mail the result to a top medical journal and, as a result, widely disseminate the cure for cancer. Here I have performed an amoral action. I have performed an action that is not morally requirable given my cognitive limitations. However, it's also the case that I have behaved in a morally permissible way: I haven't physically assault the person next to me. Hence simply because one can be morally required to perform a particular action or set of actions does not mean that one cannot also perform an *amoral* action in the effort.³⁰ Sometimes performing an action that lacks moral valence entails also performing an action that is morally permissible.

Of course, this is perfectly correct. (4) does not hold, and hence (5) also does not hold, in cases in which one can perform an amoral action coincidentally with a morally permissible action. But to deny (4) on the current grounds simply opens a new route to the conclusion. To see this, consider that there may be cases in which it is impossible to fulfill my moral obligations by performing an amoral action. Recall the original case that introduced the problem. At t , R includes: $\{\psi, \pi\}$. All one would have to do is to interpret ψ -ing, which produces far more moral good than any other members of this set and so, plausibly, is the sole member of R_p , as unable to be performed as a result of ϕ -ing, which is a non- R member of S . In this case it would appear that an *ex hypothesi* amoral action is incompatible with performing an action in R_p , and hence doing so is immoral, and hence doing so is not *amoral*.

One might respond by holding that this proposal still leaves the possibility that there could be some amoral actions, i.e., just in those cases in which one's alternatives are such that one can perform a morally permissible action *by* performing an *amoral* action. But this would seem to render certain actions amoral only in cases in which one's alternatives are such that to perform an amoral action is also to act in a morally permissible manner. This result violates IEA. Whether an action is morally evaluable for a person at a time should be a feature of the properties of *that* action—in particular, whether it displays the relevant disqualifying factors—not the extent to which that action is compatible with the performance of some *other* action. And if this is right, then if we have reason to believe that performance of some particular action ϕ can be immoral, and hence can be evaluable, for a person A in one case with one set of alternatives (such as the case of curing cancer when one is required to donate one dollar instead),

³⁰Or, if one prefers a coarse-grained scheme, an amoral action *description* in the effort.

then we have reason to believe that ϕ -ing can be morally evaluable with any set of alternatives (such as curing cancer when one is simply required not to assault one's neighbor). And hence given that to perform a purportedly amoral action is sometimes to fail to perform a morally required action, it would seem that purportedly amoral action is or can be evaluable under any circumstances. And thus not amoral.

4.2. *Immorality and Moral Requirement*

What about (6)? (6) appears to be a lynchpin in the argument: only if failing to perform a morally required action is immoral is it the case that non- R members of S are necessarily immoral, and hence not amoral. But one might deny that (6) follows from (5). One could hold, for instance, that to fail to conform to a moral requirement does not necessarily entail that one behaved in an *immoral* way. Rather, it merely implies that one behaved in an immoral way *or* an amoral way.

This is a possible response, but seems to me implausible. To deny that (6) follows from (5) requires a rather radical rethinking of the nature of immorality, in particular, the relationship between moral requirement and moral prohibition. It just seems right to say that if morality commands you to ϕ , any action other than ϕ is *immoral*, morally negative, morally *worse* than the alternative required action. One might say that an action is immoral, morally negative, morally *worse* insofar as it is both an instance of a failure to conform to a moral requirement *and* is morally evaluable. But I think this should be regarded, at best, as "Plan B". If you are morally required to do something, this is tantamount to morality saying, as it were, "*do this thing!*". And if you don't *do that thing*, I find it very difficult to believe that you've not done something morally wrong, negative, worse than the required action. That's just what it means for an action to be required: you had better do it.

4.3. *The Contrastive 'Ought'*

The final option I'd like to explore is the denial of (2), or, at least, the denial of (2) understood in a way that would yield the problematic conclusion. To see what I have in mind, consider the following account of the nature of the moral 'ought':

The One-Place 'Ought': Moral obligation is a one-place predicate, identifying a morally required act.

The one-place moral ‘ought’ is the standard operator in deontic logic.³¹ Of course, it may be convenient for us to treat moral obligation as a two-place predicate in certain cases, i.e., as a relation between a particular act and a particular agent. Also, a particular act may be required of a particular agent at t , but may, once conditions have changed at t_1 , fail to be morally required. And hence it would appear that moral obligation can be a three-place predicate, i.e., between an act, an agent, and a time:

The Three-Place ‘Ought’: Moral obligation is a three-place predicate, consisting of a morally required agent, a morally required act (or disjunctive or conjunctive set of acts), and a time at which the obligation applies.

I group these accounts of moral obligation under the heading “non-contrastive ‘oughts’”. They are non-contrastive in the following sense: for any moral obligation, there is simply one required act (or set of acts). It does not say that ϕ is required *rather than* some other act ψ . It says, simply, that ϕ is required, full stop.

But one can understand the nature of moral obligation in a different way. One can accept:

The Contrastive ‘Ought’: Moral obligation is an n -place predicate, consisting of a morally required agent, a morally required act (or disjunctive or conjunctive set of acts), a contrast class, and a time at which the obligation applies.³²

According to a contrastive account of moral obligation, one is never morally required to ϕ without the specification of a set of actions *rather than which* ϕ is morally required. A moral obligation to ϕ must always be understood as a moral obligation to ϕ *rather than* some other specified set of actions the performance of which would entail a failure to conform to one’s moral obligation.³³

³¹Alexander Broadie, “The Logical Syntax of Deontic Operators” in *The Philosophical Quarterly* 32 (1982).

³²This principle is meant to identify a wide family of views, many of which will offer different accounts of its main themes.

³³For additional arguments in favor of a contrastive ‘ought’, see Aaron Sloman, “‘Ought’ and ‘Better’” in *Mind* 79 (1970); Jackson and Pargetter, “Oughts, Options, and Actualism” in *The Philosophical Review* 95 (1986); Stephen Finlay, “Oughts and Ends” in *Philosophical Studies* 143 (2009); Justin Snedegar, “Contrastive Semantics for Deontic Modals” in *Contrastivism in Philosophy*, ed. Blaauw (London: Routledge, 2012). While these proposals differ in their various details, they capture the essence of the broad-based proposal I’m advocating here.

A non-contrastive ‘ought’ is doing heavy work in the above argument to the effect that amorality is impossible. Assume that one is morally required to ϕ , which is a member of R . According to the non-contrastive account of moral obligation, to fail to ϕ —no matter what else one does—is to fail to conform to one’s moral obligations and hence to behave immorally. And hence to perform an action that is not a member of R is to fail to ϕ , and hence to behave immorally.

If we accept the contrastive ‘ought’, however, we can insist that all members of the relevant contrast class are themselves morally evaluable, i.e., members of R . If this is correct, in performing an action that is a member of S , but not a member of R (and hence not a member of R_p), I perform an action that is not an instance of a failure to conform to my moral obligations. I only have a moral obligation to perform a member of R_p *rather than* any non- R_p members of R .³⁴ And hence if I perform a non- R member of S , this action is neither here nor there with regard to any moral obligation: I have neither fulfilled nor have failed to fulfill them. If we accept the contrastive ‘ought’, we can allow that a category of amoral actions—actions that are not morally evaluable—is possible.³⁵

To put this point in another way, (2) should be replaced with:

2*. One is morally required to perform a morally permissible action rather than a non-permissible morally evaluable action.

If we do this, given that the inference relies on (2) rather than (2*), it is clear that (5) does not follow from (4), because to perform an amoral action is not to fail to conform to a moral requirement even *if* it is impossible to perform the amoral action coincidentally with a member of R_p . One fails to conform to a moral obligation only if one performs an action outside of R_p , but within R . And hence there need be no reason to reject the possibility of

³⁴Although he denies that his view is a version of a contrastive semantics, the proposal I advocate here is also compatible with Fabrizio Cariani’s “resolution semantics” (“Ought and Resolution Semantics” in *Noûs* 47 (2013)). For Cariani, “ought”-claims only evaluate actions within the context of a salient range of “alternative courses of action available” (538). And hence it is perfectly compatible with Cariani’s view to identify this range exclusive of actions that possess the disqualifying factors. This will generate the result I advocate here.

³⁵Notice that this proposal also delivers the correct result that, e.g., while the action of *curing cancer* is amoral (i.e., is not in the relevant contrast class), the action of typing random keys need not be. After all, as noted above, this action (or action-description, if one prefers a coarse-grained account) plausibly does not possess the relevant disqualifying factors, and can, in principle, be in the relevant contrast class. Thanks to an anonymous reviewer for helpfully pointing this out.

amoral actions.

5. *Is a Contrastive ‘Ought’ Ad Hoc?*

I argue that if we wish to accept the category of amorality, we should admit the possibility of a contrastive ‘ought’, at least in contexts in which the possibility of amoral action is relevant. This ‘ought’ indicates that a particular action should be performed relative to a contrast class, i.e., rather than some other set of acts. If we accept the contrastive ‘ought’, then the category of amorality seems admissible given the conceptual constraints of the nature of moral requirement.

One might argue that this proposal is *ad hoc*. However, there are independent reasons to accept the contrastive ‘ought’. Many merits, both formal and informal, of such a view have been identified in the literature, but my favorite runs as follows. We generally discuss the relative moral merits of actions that are morally impermissible (we have “contrary-to-duty” imperatives).³⁶ For instance, imagine that you have the power to either (a) do nothing, (b) destroy Great Bend, Kansas with a nuclear explosion or (c) destroy New York City with a nuclear explosion. Of course, in this case, it is only morally permissible to (a). But we can also talk about which, of (b) or (c), one is morally required to perform, assuming that one will perform a sub-optimal action. And surely it’s the case that one is morally required to blow up Great Bend *rather than* blowing up New York City. This is a perfectly sensible claim to make, and hence it would seem that the contrastive ‘ought’ plays a role in ordinary moral discussion. But if *some* ‘ought’s are contrastive, i.e., those used in comparing morally sub-optimal acts, why should it be that the *non*-contrastive ‘ought’ is in operation when discussing the morally *optimal* action? Why should we be using an entirely different predicate to discuss the moral merits of (a)? It would seem puzzling to do so, and hence I claim that there is independent reason to accept the contrastive ‘ought’.

One might object to this line of reasoning if there were costs to accepting the contrastive rather than non-contrastive ‘ought’. In particular, one might say that the *reason* we switch from a contrastive to a non-contrastive ‘ought’ in contexts in which we move from discussing morally sub-optimal to morally optimal actions is to express the general thought that the optimal action is, well, *optimal*—it must be performed. But of course this

³⁶For the classic discussion of this issue, see Chisholm, “Contrary-to-Duty Imperatives and Deontic Logic” in *Analysis* 24 (1963).

is no reason to accept the non-contrastive ‘ought’ rather than contrastive ‘ought’. The contrastive ‘ought’ interprets the notion of moral optimality in a natural way: the morally optimal act is the one that should be performed *rather than any other*, or rather than any other action that could be morally evaluated.

6. Conclusion

I have done the following things in this paper. First, I have clarified the notion of an amoral action. Amoral actions are those that simply lack moral status. Second, I have argued that it is plausible to say that such actions should exist, stemming in large measure from the problem of moral requirability. Third, I have outlined a puzzle for the possibility of amoral actions, and finally argued that this puzzle should be solved by accepting a contrastive semantics for the moral ‘ought’.

Again, I want to stress that this paper is an exploration. I believe that there are good reasons to think that amoral actions exist, and have offered some reasons to countenance their existence. But even if you are unconvinced, I hope to have shown, at the very least, what sort of conceptual tools would be required to accept the existence of amoral actions, why someone might accept their existence, and the various means by which someone might accommodate their conceptual possibility. Whether I have identified the best possible method to accommodate them, or even whether we should accommodate them all things considered, I leave open for further exploration of this would-be moral category.