

Why Be Moral?

Terence Cuneo & Russ Shafer-Landau
(based on work done jointly with John Bengson)

ABSTRACT

The Why Be Moral? question is naturally read as a request to identify the reasons for doing as morality directs. An attractive but controversial answer to this question cites categorical moral reasons as the sorts of compelling considerations that favor doing one's duty. After elucidating the notion of such a reason (§1), we argue for the existence of moral reasons (§2), and then present an argument that some of these are categorical (§3).

Morality sometimes asks us to do what is effortless and enjoyable, such as spending time with family or praising a friend for doing the right thing. At other times, it calls on us to act in ways that are inconvenient or uncomfortable. This could be something as small as making a minor apology, or something as dramatic as risking one's life. When reflecting on such possibilities, it is natural to ask one of philosophy's perennial questions: why be moral? In the context in which it is typically raised, this question is a pointed request to identify compelling considerations that favor acting as morality requires, even if one faces competing demands or has not pledged allegiance to morality in the first place.

Our aim here is to defend the claim that there are categorical moral reasons, which are just the sorts of compelling considerations required to answer the perennial question. After explicating the idea of a categorical moral reason, we'll offer an argument designed to establish the existence of such reasons.

1. Reasons: Practical, Moral, Categorical

Let's work up to our characterization of categorical moral reasons by starting with reasons, delineating the subclass of them that are moral, and then the further subclass that are categorical.

Reasons are favorers, standing in the favoring relation to whatever they favor.¹ The reasons on which we'll focus are *practical*, favoring the performance or non-performance of actions. We'll assume that reasons apply only to beings who are capable of understanding that considerations favor or disfavor actions, of at least sometimes articulating those considerations to themselves, and of conforming their behavior to their understanding. We call such beings 'rational agents,'

¹ We'll focus on favoring, ordinarily omitting disfavoring for purposes of presentational economy.

allowing that such agents can be rational in the relevant sense even if they very often behave in ways that strike us as paradigmatically irrational or unreasonable.

Importantly, we reject a promiscuous view of reasons according to which every system of norms by itself implies a correlative set of reasons to perform the actions that those norms call for. According to that picture, norms of etiquette by themselves entail etiquette reasons; norms of chess by themselves imply chess reasons; and so forth. We recognize that one can speak of there being “chess reasons.” But this is in our view just a *façon de parler*. Chess norms do not by themselves entail any reasons at all.

Indeed, the promiscuous view that every system of norms implies its own reasons suffers from three serious defects. First, it wrongly assigns reasons of compliance to normative systems that do not warrant it. Second, it needlessly cheapens the importance of reasons—if every requirement occasioned a reason of compliance, then the question of what reason there is to comply with its edicts would be pointless. But it’s not—it is often difficult and gripping and to be taken quite seriously. Third, it generates an unhelpful epicycle in addressing the deepest questions about what there is reason to do: since some systems of guidance are inane, gratuitously demanding, or immoral, we’ll naturally be led to ask whether there is reason to comply with the reasons allegedly implied by such normative systems. By denying a promiscuous view of reasons, we sidestep this unhelpful epicycle and avoid all three concerns.

Our interest here is in *moral* reasons. Not every practical reason is moral; we recognize reasons to act that are prudential, instrumental, aesthetic, and so forth. Moreover, not every reason to act morally is a moral one; there may be non-moral reasons to act morally, as when doing the morally right thing happens to be in one’s own interest. So we need a firmer grip on what a moral reason is.

To start, note a threefold distinction between (i) the favoring relation, (ii) a fact R that favors some action ϕ , and (iii) the fact that R favors ϕ . The fact in (ii) is a reason.² The fact in (iii) is not itself a reason but what we’ll call its ‘correlative reason-fact.’ According to the characterization of moral reasons we embrace, a reason qualifies as moral when and only when its correlative reason-fact is susceptible to a particular type of explanation. Specifically:

A fact R is a moral reason to ϕ if and only if [R favors ϕ] is at least partly grounded in a moral fact of the form [If R , then ϕ is M],

² While we find a fact-based metaphysics of reasons attractive, we don’t regard it as compulsory. Theorists who prefer to think of reasons as (say) propositions, states of affairs, or tropes are welcome to paraphrase accordingly. In our estimation, nothing in our discussion hinges on these metaphysical niceties.

where M is a moral property. The intuitive idea expressed by this schema is that to determine whether a reason for action is moral we look to see whether it is morally relevant—specifically, whether its connection to the action’s moral status explains why the reason favors that action.

To see this schema at work, consider the fact [The bottle contains poison]. According to our schema, this is a moral reason to keep the bottle from a child just in case its correlative reason-fact,

[That the bottle contains poison favors keeping it from a child],

is at least partly grounded in a moral fact such as

[If the bottle contains poison, then one is morally required to keep it from a child].

Suppose it is: the correlative reason-fact holds, and is at least partly grounded in that moral requirement. It would follow that [The bottle contains poison] is a moral reason to avoid giving the bottle to a child. That would be the correct result.

In order to economize expressions, we’ll often speak of something ‘giving’ a reason, where a reason R is *given by* some x just in case R ’s correlative reason-fact is grounded in x . With this stipulation in hand, we can then say: R is a moral reason to ϕ just in case R is given by a moral fact of the form [If R , then ϕ is M]—a paradigm example of which is a moral requirement (as in the poison case above).

We turn now to clarifying the nature of categorical reasons. Such reasons contrast with hypothetical ones, where this distinction hinges on the relation between reasons and contingent commitments. We understand commitments to include desires, goals, aims, intentions, preferences, wants, plans, chosen ends, cares, or projects. *Contingent* commitments are ones that it is possible for a rational agent to lack.

Our characterization of the division between hypothetical and categorical reasons centers on the notion of a reason whose correlative reason-fact depends on an agent’s contingent commitments. More specifically:

A reason R for S to ϕ is *hypothetical* =_{def} R applies to an agent S and [R favors ϕ] is at least partly grounded in S ’s having a contingent commitment.

By contrast, a reason is categorical just in case it’s not hypothetical. In other words,

A reason R for S to ϕ is *categorical* =_{def} R applies to S and [R favors ϕ] is not even partly grounded in S ’s having a contingent commitment.

The intuitive idea is that whereas hypothetical reasons are commitment-driven, categorical reasons are not.

To illustrate, suppose a gardener's reason to plant tulip bulbs in autumn is that doing so will beautify her garden. Such a reason is hypothetical: its correlative reason-fact—[That the bulbs will beautify her garden favors her planting them]—is grounded in her having a contingent interest in enjoying the beauty of her garden. A marathoner's reason to push through to the finish line may be that completing the race will fulfill her New Year's resolution. If so, this reason is also hypothetical: its correlative reason-fact—[That completing the marathon will fulfill her New Year's resolution favors completing the race]—holds in virtue of her having the intention to fulfill that resolution. In each case, then, the reason is given by the agent's having the contingent commitments she does.

Now think of the practical reasons to save yourself from a burning building. One of these reasons may be instrumental, given by your having certain contingent commitments (e.g., to witness tomorrow's sunrise). But perhaps there is also a prudential reason for you to save yourself, one that isn't given by any of your contingent commitments. This is a candidate non-moral example of a categorical reason.

The categorical reasons we're most interested in, of course, are moral ones. The question is whether any such reasons exist. The project in the remainder of this essay is to justify an affirmative answer, and so to defend the thesis we'll call

Categoricity: There are categorical moral reasons.

While the Categoricity thesis maintains that some moral reasons are categorical, it is neutral on whether any moral reasons are hypothetical. This thesis is also compatible with both affirmations and denials of the claim that rational agents have *non-contingent* commitments that give at least some of their moral reasons. It may be easier to vindicate the existence of categorical moral reasons if we could help ourselves to this assumption. For if there were such commitments, then they might give moral reasons for all rational agents to do as morality directs. While nothing we say is intended to cast doubt on this possibility, we'll henceforth assume that all commitments are contingent. Doing so ensures that we do not make things too easy on ourselves by assigning to rational agents a commitment whose fulfillment requires moral conduct.

2. Defending Moral Reasons

There are categorical moral reasons only if there are moral reasons. It's not obvious that there are.

After all, as we've indicated, a moral reason isn't just any reason that favors doing what morality calls for—rather, a moral reason is one that is given by a moral fact (i.e., its correlative reason fact is grounded in a moral fact). Some prominent views of reasons imply that no such reasons exist. For example, instrumentalism and rational egoism are committed to the claim that all practical reasons are given exclusively by facts about how to get what you want, or how to promote or protect self-interest, respectively. These are not moral facts. It follows that those views cannot recognize the existence of moral reasons.

We dedicate this section to defending the claim that there are moral reasons. This work lays the foundation for our case in the next section that at least some of these reasons are categorical.

We offer the following argument for moral reasons. Its first premise states:

1. If an agent is blameworthy for his actions, then there are reasons for him to have refrained from those actions.

We understand an agent to be blameworthy just in case the agent merits the reactive attitudes of censure, disapprobation, or condemnation.³ That understanding lies behind the reference to blameworthiness in the second premise as well:

2. Some agents are blameworthy for their immoral actions.

It follows from these two premises that

3. There are reasons for those agents to have refrained from their immoral actions.

Now for two additional premises:

4. At least some of these reasons are given by a moral requirement.⁴
5. If a reason is given by a moral requirement, then the reason is a moral one.

³ Being blameworthy isn't merely to be liable to merited criticism, since such criticism can fall short of blame. Note that when an agent is blameworthy, this does not imply that anyone ought to *convey* censure, disapprobation, or condemnation; there may sometimes be compelling reasons to refrain from expressing merited blame. And there may be cases in which it is justified, all things considered, to express blame even when the target of opprobrium is not blameworthy.

⁴ We focus here and below on requirements (rather than prohibitions) just for ease of exposition; all points made in terms of requirements can be made in terms of prohibitions, *salve veritate*.

So,

6. There are moral reasons.

Since the claim in 5 is a direct implication of the understanding of moral reasons we provided in §1, we won't pause to discuss this premise here. Instead, we'll focus on the argument's other three premises in the order they're presented.

Premise 1 expresses the general principle that blameworthiness presupposes reasons. If there are no reasons for you to refrain from performing a given action, then performing it cannot merit blame. Being deserving of blame entails that you have ignored, disregarded, or otherwise neglected a relevant consideration that opposes the action you have performed. If you have properly complied with all relevant reasons—if you have appreciated all such reasons, given them their due weight, and acted accordingly—then you are not blameworthy for having performed the action. So an agent is blameworthy for an action only if the agent has failed to comply with some reason that stands against it. That is what the first premise says.

Although this defense of premise 1 may seem fairly simple, the issues it covers are in fact both significant and complex. So let us dig deeper, probing what are arguably the most pressing objections that can be leveled against it. We consider three such objections here.

The first appeals to putative counterexamples, such as the following: a person intentionally sets out to do something that is just awful, morally speaking. He has no justification and no excuse. So he is blameworthy. But through his inadvertence, or some fortunate accident, what he ends up doing isn't at all bad, and may even be quite good. For example, imagine a disgruntled employee who believes that the coffee he is about to hand a co-worker is laced with poison. Unbeknownst to our wrongdoer, the "poison" is just saccharine, and the coffee is served up exactly to his co-worker's tastes. Surely such a person is blameworthy for what he's done. Yet there is no reason to avoid handing someone a delicious cup of coffee. So we have warranted blame for an action without there being a reason to have acted otherwise. Thus premise 1 is false.

We agree that, all else equal, there is no reason to refrain from handing a perfectly safe cup of coffee to someone who will enjoy it. But that is not what we are blaming the employee for. We are blaming him for attempting to poison his co-worker. And there is a reason not to do that. That the employee ignores this reason, or treats it as favoring his conduct, explains why he is culpable for his action. What this case illustrates is that, when an agent is blameworthy, we must attend carefully to the specific aspect of his behavior that merits criticism. The mere fact that there are descriptions under which an agent's actions are innocuous

or even attractive is insufficient to establish that he is not, in virtue of performing those actions, flouting a reason to act otherwise.

A second route to criticizing premise 1 claims that blameworthiness does not entail the flouting of reasons, but rather just the presence of immoral motivations or intentions. The idea here is that agents are blameworthy for their misdeeds because, in performing them, such agents exhibit various vices. Thus vicious actions warrant blame even in cases where the agents lack reasons to act otherwise than they do.

We demur. Suppose, for instance, that a jealous, spiteful person performed his moral duty. If we know nothing more about the case, then, on this basis alone, we have no reason to regard him as blameworthy for his conduct. This is not to immunize him from other sorts of criticism. Indeed, we may appropriately issue a negative assessment of his character, various aspects of which could even merit blame. Further, we might also rightly blame him for having discharged his duty resentfully, snidely, or begrudgingly, though such blame would be merited only if there is a reason not to behave in such a manner. No matter how vicious, a person is not deserving of blame for performing an action, or for performing it in a particular manner, if there are no reasons that oppose his doing so.

Turn now to the third objection to premise 1. This objection says that agents are blameworthy for actions if they violate a norm of conduct in good-standing—even if they lack a reason to do what that norm says. The objection assumes that such a norm need not give reasons of compliance. Rather, the norm is justified on other grounds. For example, general conformity to it would achieve a valuable social goal, such as coordinating behavioral expectations, encouraging agents to act in non-harmful ways, or signaling to community members that certain actions are apt for sanction. The core thought here is that warranted blame occurs within a practice that is defined by a set of rules that are in good standing (such as those that are effective in achieving a valuable social goal); blame is deserved whenever these rules call for it, and they may do so even when there is no reason for a given individual to abide by these rules. As a result, premise 1 is false.

While we agree that warranted blame occurs within a practice, we find the line of thought underlying this objection unpersuasive. Specifically, we deny that mere failure to conform to a norm in good standing thereby renders an agent blameworthy. Although the norm to observe public debates between candidates for public office is in good standing, since it promotes a valuable social goal (such as ensuring an informed electorate), a given agent is not blameworthy for flouting this norm if there is no reason for that agent to conform to it.

Merited blame points to a shortcoming. When the object of such blame is not a collective but an individual, as in our cases, that shortcoming implicates the

agent. Although she may be liable to certain sorts of criticism for flouting norms in good standing, she does not merit the reactive attitudes of censure, disapprobation, or condemnation for flouting a norm that fails to give her any reason for obedience. Those attitudes would be deeply unfitting. What is more, any norms that license blame must allow for the existence of legitimate excuses; an agent has such an excuse if there are no reasons for her to comply with those norms. The objection we are considering severs the tie between blameworthiness and personal shortcoming, sanctions unfitting reactions, and fails to recognize a class of legitimate excuses. Premise 1 preserves the connection between blameworthiness and personal shortcoming, bars the relevant unfitting reactions, and can help explain why agents are sometimes blameless for the actions they perform. That is excellent reason to accept this premise and to reject the third objection.

Consider next the argument's second premise, which states that some agents are blameworthy for having acted immorally. We regard this premise as part of common sense. It has of course been denied. But such denials do not occupy a privileged position that would justify either the rejection of our view or the disruption of our efforts to defend it. What is more, the premise can be defended by observing that agents sometimes inexcusably perform wrongful actions on the basis of uncoerced choices that are expressive of their character. It is fair to say that if an agent acts in this way, then he is blameworthy for his action. Some agents have indeed acted in that way. It follows that some moral agents are blameworthy for their immoral conduct, just as premise 2 says.

Turn now to premise 4, which tells us that at least some of the reasons are given by moral requirements. We can defend this premise by identifying examples in which agents flout reasons that apply to them, are blameworthy, and whose blameworthiness is explained by their failure to conform to moral requirements. This confluence, we contend, is best explained by the claim that those reasons are given by these requirements.

Consider a case in which you have promised to provide a meal for an injured colleague and her family. Suppose that, rather than keeping your promise, you fritter the day away by playing video games. You would be blameworthy for doing so. We know from premise 1 that there must be a reason for you to have acted otherwise, since you are blameworthy. Furthermore, your blameworthiness is explained by the fact that you are morally required to provide a meal and yet failed to do so—that's what you are blameworthy *for*.

What could explain this confluence? Suppose the reason to fulfill your promise is given by the moral requirement to provide a meal. That is, suppose the reason's correlative reason-fact,

[That you've promised to prepare a meal favors providing it],

is grounded in the fact

[If you promised to prepare a meal, then you're required to provide one].

That would make good sense of things. By contrast, were the reason and requirement not related in this way, the fact that you are blameworthy for violating the requirement, by failing to do what you have reason to do, would be utterly perplexing.

This line of reasoning easily generalizes to other moral requirements and prohibitions. Consider, for instance, the requirement to protect your children from lethal danger, or a prohibition against engaging in recreational slaughter. Absent excuse, those who violate these demands are blameworthy, which entails that there is a reason for them to have acted otherwise. And their blameworthiness is explained by the fact that they have flouted the requirement or prohibition in question. In any such case, and in a great many others we could easily think of, this confluence is best explained by viewing the requirement or prohibition as giving the reason. That is just what premise 4 says.

This concludes our defense of the existence of moral reasons. We turn now to arguing for the claim that some of these are categorical.

3. Defending the Categorical Thesis

Our argument for the claim that some moral reasons are categorical highlights a figure we'll call the 'immoralist.' The immoralist is an agent who is fully informed about non-normative matters and has a perfectly consistent set of commitments. He is highly capable of assessing options, assiduously gathers information to discover how best to pursue his chosen goals, and takes the needed steps to ensure that his goals are met. He is also, crucially, entirely indifferent or opposed to standard moral demands, and has no contingent commitments that would be served by toeing the moral line. The character we are imagining has not been brainwashed or neurologically manipulated into having his odious psychological profile, but autonomously aligns himself with evil ends. While such a person needn't always choose evil ends or deeds as such (i.e., *de dicto*, under the guise of evil), he does delight in such things as humiliating the weak, breaking promises that others heavily rely on, or, at the limit, killing people for sport. He forsakes opportunities to easily rescue others from dire situations and takes immense pleasure in watching the fatal consequences unfold. He discriminates against others based on their appearance. He attaches no intrinsic importance to the interests of others. And so on.

We focus on the figure of the immoralist not because his misdeeds are

uniquely disfavored. As we shall see, that position is untenable. The immoralist simply serves as a heuristic device in moral philosophy, in much the same way as a frictionless plane does in physics. In particular, the figure of the immoralist enables us to clearly formulate our argument for the Categoricality thesis, which has two premises:

A. If there are reasons for the immoralist to refrain from his immoral actions, then at least some of those reasons are categorical moral reasons.

B. There are reasons for the immoralist to refrain from his immoral actions.

So,

Categoricality: There are categorical moral reasons.

We will try to reveal the attractions of both premises of this argument in what follows.

3.1 Immoralists and their reasons

Premise A of the argument is meant to be acceptable to both fans and critics of categorical reasons. The immoralist we are imagining is flouting widely acknowledged moral requirements, and is precisely a person who lacks any contingent commitments that would be furthered were he to hew to the moral path. So *if* there are any reasons for him to refrain from his evil deeds, then those reasons are categorical ones, given by the moral requirements that direct him to so refrain. This constitutes our defense of premise A.

Critics might push back in one of two ways. A first objection insists that the immoralist must have *some* contingent commitments that would be served were he to refrain from his wicked ways. On such a line, whatever reasons the immoralist has to desist will stem straightforwardly from his aversion to jail, or his desire to avoid the potential harms inflicted by his vengeful victims, and so on. That would enable these critics to resist premise A, since the immoralist's reasons would then be hypothetical, rather than categorical.

However, this response is unavailable, since we have stipulated that the immoralist does *not* have such commitments. In reply, the critic is likely to object that this makes the immoralist we have characterized a mere fantasy, for all people in the real world have some commitments that will be served were they to refrain from moral wrongdoing. Yet even if all actual people do have some contingent ends that would be served by avoiding cruelty, we can specify a possible world in

which there is an immoralist who does not.⁵ And that is all that is needed to establish the claim that the reasons to which premise A refers are categorical.

To further allay the worry that the immoralist is purely fictitious, imagine someone very much like our immoralist, but who has, like the rest of us, some contingent commitments that *do* give him hypothetical reasons to refrain from immoral behavior. Such an agent might, for example, have a strong aversion to being found out and punished for his misdeeds. As we see it, this aversion might give him an additional reason not to inflict misery on others, one that is likelier than a moral reason to motivate him to do the right thing. However, such additional reasons—whether hypothetical or not—are not the whole story. To appreciate this, notice that the cruelties he intends to perpetrate are prohibited by norms that make no mention of his aims or well-being. These norms will, first and foremost, cite the suffering of his victims, his disregard for their welfare, and the absence of their consent to his treatment. If there are reasons that oppose his misdeeds, at least some of them will be given by these norms, which are not commitment-driven. Such reasons would not be hypothetical but rather categorical.

From here it's a small step to establishing that the reasons would be *moral*, as premise A asserts. After all, the norms just mentioned are moral requirements or prohibitions that direct agents to refrain from such actions. So if the reasons are given by those norms, then (per our definition of moral reasons in §1) those reasons are moral ones.

Supposing we are right about this, then premise A is secure: *if* there are any reasons that oppose the immoralist's actions, then there are categorical moral reasons. It's time to consider whether the antecedent of that conditional—premise B—is true.

3.2 *The blameworthiness of immoralists*

Premise B says that there are reasons for the immoralist not to do what he has his heart set on. We can summarize our defense of this premise thus: Agents who perpetrate serious immoralities, while possessed of deeply immoral motivations and intentions, and lacking in any excuse, are blameworthy for their deeds. The immoralist we are envisioning meets all of these conditions. His actions are as bad as can be. His motivations and intentions are no better. And he is a rational agent, responsible for his deeds. It may be a gross understatement to say so, but when such an agent succeeds in his aim, he is blameworthy for his actions.

⁵ Many fantasies depict genuine possibilities. Sharon Street (2009, §4) agrees that immoralists are possible but claims that when we are sufficiently imaginative to picture what such a person is like, we'll no longer be tempted to think that there are reasons for him to refrain from such cruel treatment. But Street offers no argument for this claim, focusing instead on raising objections to the different claim that we have certain *prudential* reasons that are not grounded in contingent commitments.

As we argued above, however, one is blameworthy for an action only if there is some reason to refrain from committing it. Because the immoralist is blameworthy for his deeds, there is a reason that opposes his actions. That he has violated or ignored it is the basis of his blameworthiness.

We have here the makings of an argument for B. Its first premise is identical with premise 1 in our argument for the existence of moral reasons, repeated here:

1. If an agent is blameworthy for his actions, then there are reasons for him to have refrained from those actions.

The argument's other premise is a variation on premise 2 in the argument for moral reasons:

- 2*. The immoralist is blameworthy for his immoral actions.

These two premises entail:

- B. There are reasons for the immoralist to have refrained from his immoral actions.

We believe that each of the premises in this argument is highly plausible. We have defended the first premise in §2, and so focus our efforts here entirely on the second premise, 2*.

This premise tells us that immoralists are blameworthy for their terrible deeds. It can be supported thus: If *any* agents are blameworthy for their actions, surely those who are voluntarily bent on the gravest evil are among them. This is so whether immoralists are doing evil for its own sake, or doing what is in fact evil, all the while characterizing their actions to themselves as aimed at a good. An informed, uncoerced immoralist is the perfect exemplar of the blameworthy agent. His consistency is no proof against criticism. His intelligence and cunning, his ability to select appropriate means to his chosen ends, and the meticulous care with which he prepares to hurt others render him more, rather than less, liable to blame.

We can defend this verdict, and so 2*, by means of the same line of reasoning that we used in our defense of premise 2 (in §2). There, we maintained that if an agent inexcusably performs an action that is morally wrong on the basis of an uncoerced choice expressive of his character, then he is blameworthy for his action. Now consider the immoralist. He has performed morally terrible deeds. Moreover, he's done so on the basis of uncoerced choices that are expressive of his character. Furthermore, he lacks an excuse for his characteristic wrongdoing. These three points jointly support the claim that, and explain why, he is blameworthy for

his terrible deeds.

The first two points are simply part of our stipulation regarding the figure of the immoralist. The third point follows from the observation that the standard conditions for exculpation do not apply here. The dedicated evildoer we are imagining is not compelled to act as he does, but has chosen his path and has ruthlessly pursued it in the absence of duress, coercion, or factual ignorance. His terrible deeds are as inexcusable as an agent's actions can be.

Note that the argument we've offered is compatible with the possibility that there is a moral monster of some description whose wrongdoing *is* fully excused. Perhaps, for instance, such a person was brainwashed from an early age by members of his community, prevented from leaving the community, and raised in such ignorance as to be unable to engage in the sorts of reasoning that might lead him to the moral truth. These conditions might together fully exculpate this person—call him the 'incurable evildoer.' But even if the conditions were exculpatory, that would in no way threaten the case we've made for categorical moral reasons. We offer two observations to support this verdict.

First, our argument requires just a single example in which the figure introduced above is blameworthy for his evil deeds. The original case we described satisfies this description. The immoralist is a paradigm of blameworthiness, even if his incorrigible cousin is not.

Second, even if incorrigible evildoing is an excuse,⁶ this gives us no reason to suppose that anything less than incorrigibility is exculpatory. If the incorrigible evildoer is to be excused for his bad behavior (a matter about which we remain agnostic), that is presumably because there was *never* a moment in which he was able to autonomously elect good over evil. Here in the real world, however, such moments are available to nearly anyone—and, probably, to everyone. So it's likely that everyone, or perhaps nearly everyone, is subject to categorical moral reasons.

One might think that the incorrigible evildoer poses a problem for our view for a different reason, expressed by the following argument: If there is categorical moral reason for one rational agent to ϕ , then there is categorical moral reason for every rational agent to ϕ . But for every morally required action, there will be at least one such agent—the incorrigible evildoer—who lacks such a reason. So there are no categorical moral reasons after all.

Both elements of this line of reasoning are mistaken. The initial conditional is false; there are many categorical moral reasons that apply to only some, not all, agents. There is, for example, a categorical moral reason for those who issue racist comments to apologize for their conduct, a reason that their victims lack; that you

⁶ Wolf (1987) maintains that such incorrigibility is an excuse. *Cp.* Mason (2019, ch. 6), whose view implies that it excuses agents from what she calls "ordinary blame," though not from so-called "detached blame" (which is directed to agents who lack commitment to objectively correct moral standards).

promised to help your neighbors gives you a categorical moral reason that most others lack (since they've made no such promise). Further, in our view, the assumption that there is no categorical moral reason for the incorrigible evildoer to refrain from his awful conduct is false. To be clear, we are allowing that he is blameless, if his conduct is indeed fully excusable. However, that he is fully excusable (if he is) is one thing; that there was no reason opposing his conduct is quite another. An excuse for bad behavior does not extinguish the reasons that oppose it; in fact, excuses for poor conduct presuppose the existence of reasons not to engage in it. Arguably, Jean Valjean should be excused for stealing bread to feed his sister's children. Still, there is some reason for him not to steal. If there weren't, then his actions wouldn't be apt for exculpation in the first place.

While the line of thought we've pursued in our argument for the existence of categorical moral reasons has taken a variety of twists and turns, we can succinctly summarize it as follows. Premise 1 of the argument says that if there are reasons for the immoralist to act otherwise than he does, then at least some of them are categorical moral reasons. We defended this premise in §3.1, by drawing attention to our stipulations regarding the immoralist's choices, character, and aims. Premise 2 says that there are indeed reasons for the immoralist to act otherwise than he does. We defended that premise in this subsection, by means of an argument whose first premise, borrowed from our argument in §2, registers a general connection between blameworthiness and reasons; its second premise—that the immoralist is blameworthy—is a direct consequence of his having inexcusably performed morally terrible deeds on the basis of uncoerced choices that are expressive of his character.

Conclusion

If our arguments have met with success, they show that there are categorical moral reasons to do as morality directs. Such reasons are well-suited to answer the Why Be Moral? question, which, on a natural reading, is a request to identify compelling considerations that favor moral action, even when it comes at the cost of frustrated desires or sacrifices of self-interest. Categorical moral reasons don't, of course, literally compel—no reasons can do that. But normative authority is not a matter of compulsion. It is rather a matter of being backed by genuine reasons that obtain even when one is disinclined to honor them. Categorical moral reasons are authoritative in just this way.