*Episteme* (2025), 1–14
doi:10.1017/epi.2025.10056

episteme

ARTICLE

# When it is (and is not) Blameworthy to Break the Rules

Vishnu Sridharan

Assistant Professor, University of Colorado Boulder, Affiliate Faculty, Colorado Law, Boulder, USA
Email: vishnu.sridharan@colorado.edu

(Received 28 December 2024; revised 29 March 2025; accepted 18 May 2025)

**Abstract**

While the claim that moral ignorance exculpates is quite controversial, the parallel claim with respect to non-moral ignorance seems to be universally accepted. As a starting point, we can state this claim as follows:

Non-moral Ignorance Exculpates: If an agent did everything that could be reasonably expected of her to inquire into some empirical issue as to whether P, the seeming truth of P played the appropriate role in the agent's motivation to Φ, and the agent would not have merited blame for Φ-ing if P had been the case, then the agent does not merit blame for Φ-ing.

In this paper, I aim to accomplish two tasks. First, I argue that NMIE is false in certain cases in which, by Φ-ing, the agent violates a course-grained, reasonable community norm without knowing that doing so is in everyone's best interests. Second, I argue that, while moral ignorance, like non-moral ignorance, does not exculpate when community norms are violated in this manner, it does exculpate when they are not. With these two tasks accomplished, we will see the striking parallels in the manner in which both moral and non-moral ignorance exculpate.

**Keywords:** Moral ignorance; knowledge; blame; norms; risk

While the claim that moral ignorance exculpates is quite controversial, the parallel claim with respect to non-moral ignorance seems to be universally accepted.[1] As a starting point, we can state this claim as follows:

Non-moral Ignorance Exculpates (NMIE): If an agent did everything that could be reasonably expected of her to inquire into some empirical issue as to whether P, the seeming truth of P played the appropriate role in the agent's motivation to Φ, and the agent would not have merited blame for Φ-ing if P had been the case, then the agent does not merit blame for Φ-ing.

[1]At least arguably, this intuition goes back to Aristotle (2009), who discussed the manner in which actions that stem from ignorance are "involuntary" in Book III of the Nicomachean Ethics. A more recent defense of this claim can be found in Zimmerman (1997) and, to some extent, Smith (1983). I have come across no counterexamples to the claim that non-moral ignorance exculpates in this manner.

As an example, if YiFei did everything that could be reasonably expected of her to inquire into whether the coffee she is serving is safe for consumption, the fact that the coffee seemed safe played the appropriate role in her decision to serve it, and she would not have merited blame for serving the coffee if the coffee were safe for consumption, then she does not merit blame for serving coffee, even if that coffee is poisoned. Otherwise put, as long as YiFei has done all that she ought to have done to figure out whether the coffee is safe to consume, then whether or not the coffee actually is safe ought not to contribute to her blameworthiness for serving it.[2]

In this paper, I aim to accomplish two basic tasks. First, I argue that, in a range of important cases, NMIE is false. More specifically, NMIE is false in certain cases in which, by Φ-ing, the agent violates a course-grained, reasonable community norm without knowing that doing so is in everyone's best interests.[3] Second, I argue that, while moral ignorance, like non-moral ignorance, does not exculpate when community norms are violated in this manner, it does exculpate when they are not. With these two tasks accomplished, we will be able to see a parallel in the manner in which both moral and non-moral ignorance can exculpate. In particular, we will see that both moral and non-moral ignorance can be exculpated, at least in a range of ranges, with knowledge that one's action is in everyone's best interests.

In the context of non-moral ignorance, it is generally agreed that ignorance can exculpate when an agent has done everything we can reasonably expect of her to inquire into some empirical matter. Most people are amenable to the thought that, when an agent has fulfilled these epistemic obligations, ignorance of some relevant proposition (or set of propositions) may be exculpatory. What is likely to require more argument is that there are cases in which such ignorance cannot exculpate: convincing the reader of this is the task of Section 1.

The context of moral ignorance provides us with the opposite challenge, since few people are amenable to the thought that an agent's epistemic limitations can render her blameless when she performs actions that are morally wrong. To respond to this challenge, in Section 2 below, I first show that the primary cases that have been discussed in the literature on moral ignorance are ones in which agents violate coarse-grained, reasonable community norms without knowing that doing so is in everyone's best interests. Since, in my view, even non-moral ignorance does not exculpate in cases like these, we should be unsurprised that moral ignorance does not either. However, in cases in which an agent does not violate norms in this manner, there is a much stronger case to be made that moral ignorance does, in fact, exculpate. If we are tempted to exculpate ignorant agents when they do not violate norms in this manner—and in particular in certain cases when they know that their actions are in the best interests of all—then we will see significant parallels between the cases in which both moral and non-moral ignorance exculpate.

## 1. Non-moral ignorance and exculpation

### *1.1. The basic claim*

Despite our best efforts, we often fail to learn crucial facts about the world around us. In fact, when it comes to almost any proposition about the external world (and perhaps the

[2]It's worth noting that even this case of non-moral ignorance may involve some amount of moral ignorance. For instance, YiFei may believe that serving the coffee is not wrong even though, given the fact that it is unsafe for consumption, it actually is.

[3]We may also frame this in terms of rights violations. For instance, Quong (2020: 200) talks of gambling with the rights of others.

https://doi.org/10.1017/epi.2025.10056 Published online by Cambridge University Press

internal world, for that matter), no amount of effort and good intentions can ensure that our actions will be well-informed. To continue with the example discussed above, YiFei might do absolutely everything we might expect of her to figure out whether the coffee she is serving is safe but still end up serving unsafe coffee. With this in mind, it's not particularly controversial that, in a wide range of contexts such as YiFei's, non-moral ignorance can exculpate.

As stated above, the basic formulation of the claim that non-moral ignorance exculpates that I will work with is as follows:

> Non-Moral Ignorance Exculpates (NMIE): If an agent did everything that could be reasonably expected of her to inquire into some empirical issue as to whether P, the seeming truth of P played the appropriate role in the agent's motivation to Φ, and the agent would not have merited blame for Φ-ing if P had been the case, then the agent does not merit blame for Φ-ing.

Before examining counterexamples to NMIE, it's worth saying a bit about each of its conditions. While there may be disagreement about what can be reasonably expected an agent do to inquire into whether some proposition is true, potential activities might include seeking the guidance of experts, reading relevant literature, and perhaps even checking one's reasoning against that of one's peers. In the counterexamples to follow, the exact nature of our expectations will not be crucial, as long as we agree that an agent might do everything expected of her to inquire into the truth of a proposition but still end up lacking crucial information.

The purpose of NMIE's second condition is to screen out cases in which the agent would have Φ-d regardless of what she learned through her inquiry. For instance, if YiFei would have served the coffee regardless of what she learned through her research into the coffee's safety, then she would likely merit blame for serving the coffee even if it was not poisoned. The purpose of NMIE's third condition is to screen out cases in which the agent might merit blame for Φ-ing for reasons that may be unrelated to whether the proposition in question is true and that proposition's role in the agent's motivation. For instance, the agent might merit blame for Φ-ing because she had sinister motivations, such as the case in which YiFei served coffee with the hope that it would prevent a colleague she disliked from getting any sleep. In such a case, YiFei might merit blame for serving the coffee even if she did everything we could reasonably expect of her to inquire into its safety and the seeming safety of the coffee played the appropriate role in her motivation to serve it.

### *1.2. Counterexamples*

With this basic understanding of NMIE, we can take a look at some potential counterexamples. To start, consider the following:

> FROZEN BRIDGE: Bus driver Betty has encountered some extremely dangerous road conditions. With a bus full of passengers, she has two choices: the 'Prudent Route' or the 'Risky Route.' If Betty takes the Prudent Route, all of her passengers will sustain moderate but non-fatal injuries. If Betty takes the Risky Route and the bridge has not frozen over, all of her passengers will be unharmed. However, if Betty takes the Risky Route and the bridge has frozen over, all of her passengers will die. As a bus driver, Betty knows that she ought to be cautious when other people's lives are in her hands. Betty investigates the bridge conditions to the best of her ability; unfortunately, despite her best efforts, Betty forms the false belief that the

https://doi.org/10.1017/epi.2025.10056 Published online by Cambridge University Press

bridge has not frozen over. Guided by this false belief, Betty takes the Risky Route and all of her passengers die.

In FROZEN BRIDGE, Betty seems to merit blame for taking the Risky Route. In short, this is because Betty violates a coarse-grained, reasonable community norm without knowing that doing so is best for everyone.[4] The coarse-grained[5], reasonable community norm, in this instance, would be something like "When other people's lives are in your hands, play it safe." This simple formulation would capture something like the minimax principle, or the belief that we ought to minimize the largest complaint that others will have against our course of action.[6] While all violations of norms are not blameworthy, when Betty violated this norm she did not know that doing so would be best for everyone. Otherwise put, even though taking the Risky Route had the potential for a huge moral payout—none of her passengers being injured—it was both more prudent for her to take the Prudent Route and she didn't know that she would do better otherwise. For this, Betty merits blame.

We can contrast FROZEN BRIDGE with the following sort of case:

SAFE BRIDGE: Like Betty, bus driver Angela has encountered extremely dangerous road conditions and faces a choice between a Prudent Route and a Risky Route. Angela also knows that she ought to be cautious when other people's lives are in her hands. Angela investigates the bridge conditions to the best of her ability and, through her research, comes to *know* that the bridge has not frozen over.[7] Guided by this knowledge, Angela takes the Risky Route and none of the passengers are harmed.

Since Angela knew that, if she took the Risky Route, none of the passengers would be harmed, Angela does not seem to merit blame for her action. After all, if she took the Prudent Route, all of her passengers would have sustained moderate injuries. Since Angela knew that taking the Risky Route, in her case, did not actually pose a risk to any of her passengers, there doesn't seem to be any reason to criticize her for her course of action.[8] Otherwise put, Angela knew that her course of action was best for everyone, so her violation of a coarse-grained, reasonable community norm does not seem blameworthy.

One objection we might have to the idea that Betty merits blame in FROZEN BRIDGE but Angela does not in SAFE BRIDGE is that, in order to accept this, we must accept that whether one merits blame can depend on whether one knows, as opposed to falsely believes, that one's course of action is morally best. Angela knows that taking the Risky Route is morally best in virtue of knowing that no one will be harmed as a result.

[4]Perhaps the most thoroughly developed account of such reasonable community norms comes from Scanlon's contractualism (see especially 1998). I focus on coarse-grained, as opposed to more finely grained, community norms because those are ones that we can most easily expect others in our moral community to have some awareness of. I revisit this point below. As Scanlon (1998, 205) writes on this point:

There is an obvious pressure toward making principles more fine-grained, to take account of more and more specific variations in needs and circumstances. But there is also counterpressure arising from the fact that finer-grained principles will create more uncertainty and require those in other positions to gather more information in order to know what a principle gives to and requires of them.

[5]The importance of fineness of grain is explored in more detail below on p. 19.

[6]Of course, my purpose is not to argue for this principle as a reasonable community norm. Regardless of what we take to be reasonably community norms, we can formulate cases like FROZEN BRIDGE.

[7]If the reader thinks it matters, Angela can also know that she knows that the bridge has not frozen over.

[8]I say more about the relevant notion of 'risk' below.

https://doi.org/10.1017/epi.2025.10056 Published online by Cambridge University Press

Betty falsely believes that taking the Risky Route is morally best, but the gap between her false belief and knowledge results in all of the passengers being killed. If we adopt the notion of blame being put to use here, then, knowledge can play a significant role in our moral evaluations of other's behavior.

While I think it's clear that knowledge has an important role to play in moral evaluations, many may hesitate before evaluating Angela any differently than Betty. In particular, they might argue that, if Betty failed to do what could be reasonably expected of her, then surely Angela failed to do what could be reasonably expected of her as well. In response, we should be clear that, while Angela and Betty are superficially similar agents, the differences between them are stark. In Betty's mind, taking the Risky Route had the potential for a huge moral payout; Angela, on the other hand, knew that taking the Risky Route would achieve this payout. In other words, while Betty was putting the lives of her passengers at risk, Angela knew both that her passengers had nothing to worry about and that, by taking the Risky Route, each passenger would emerge in a better position.[9]

At this point it's worth pausing to say something about risk. While my primary interest is in a subjective notion of risk—or subjective probabilities—objection notions illustrate a distinction between Angela and Betty as well. (In particular, the objective risk that the fridge is frozen in Angela's case is zero, while the objective risk that the bridge is frozen in Betty's case is 1). One assumption I will make to start is that the propositions that an agent knows are part of her evidence.[10] With this in mind, it will be part of Angela's evidence that the bridge is not frozen, and that taking the Risky Route will not result in harm to any of her passengers. Since these propositions are entailed by Angela's evidence, there is no subjective risk that either one is false. In other words, I take it to be at least defensible that Angela can be fully confident in the propositions that she knows[11], even if, in certain high-stakes situations, she must meet a higher epistemic standard to achieve such knowledge.[12]

Of course, I need one additional assumption in order to draw a distinction between the risk that Angela exposes her passengers to—which I argue is none—and the risk that Betty exposes her passengers to. This additional assumption is that an agent who knows a proposition has stronger evidence for that proposition than an agent who simply falsely believes it. In other worse, the epistemic position of an agent who knows $p$ is stronger, vis a vis the proposition p, than an agent who falsely believes it. Even if we think Betty did all that she could to learn whether the bridge was frozen, I don't think we should be committed to the claim that her epistemic position with regard to that proposition is just as strong as Angela's. Instead, perhaps through no fault of her own, Betty failed to achieve knowledge while Angela succeeded in doing so.

If Angela's evidence entails that the bridge is not frozen, and her evidence is stronger than Betty's, then Betty will be exposing her passengers to a risk that Angela is not. Perhaps more simply put, when we know that our actions will leave everyone better off, we do not expose them to risk in the same way as we do when we falsely believe this. In this way, even a notion of risk that is centered around subjective probability can accommodate the importance of knowledge to our moral evaluations.

[9]For further discussion, see Sridharan (2020, 2024).

[10]This is not equivalent to E = K, as defended by Williamson (2000), since I remain open as to what other propositions (or phenomenal states) might also be part of an agent's evidence.

[11]If we don't think full confidence is justified here, then we can simply draw a line between the confidence merited by knowledge that merited by false belief.

[12]On so-called pragmatic encroachment, see, for instance, Brown (2008) and Anderson and Hawthorne (2019).

https://doi.org/10.1017/epi.2025.10056 Published online by Cambridge University Press

Before moving on, one additional objection is worth addressing. According to this objection, whatever norm Betty is guilty of violating in FROZEN BRIDGE is one that Angela also violates in SAFE BRIDGE. If Angela violates the same norm as Betty, this objector continues, then surely they are equally blameworthy. As noted above, I agree that both Angela and Betty violate a coarse-grained, reasonable community norm. However, I think we should reject the view that it is always blameworthy to violate such norms. We simply cannot expect that rules or social norms will always capture the appropriate course of behavior in any context, especially in coarse-grained form. Rules can provide excellent guideposts as to how we ought to comport ourselves and treat others, and they can serve to stabilize expectations around such comportment.[13] At the same time, we should acknowledge that there will be cases in which rule-breaking will not be worthy of criticism (in fact, at times it may be worth of praise). For instance, in SAFE BRIDGE, Angela knew that her course of action would be best for everyone involved. Since each and every passenger is better off for Angela taking the Risky Route—and Angela knew this would be the case ahead of time—she does not seem to merit blame for violating this coarse-grained, reasonable community norm.[14]

We'll return to these and other objections in more detail below. For now, it's simply important to establish that Betty's blameworthiness is compatible with fulfilling both of NMIE's conditions. Let's take the conditions one by one, with the proposition p being that the bridge has frozen over.

First, it seems clear that, in a case like FROZEN BRIDGE, Betty might have done everything that we could reasonably expect of her to inquire into whether the bridge had frozen. Betty may have talked to a number of climate and engineering experts, and she might have taken a close, hard look at the digital and analog data available to her. Regardless of the lengths to which we think Betty ought to go to in her inquiry, there seems to be no principled reason why she could not do so in Frozen Bridge. In addition, the fact that it seemed like the bridge had not frozen over can easily have played the appropriate role in Betty's motivation to take the Risky Route.

In order for FROZEN BRIDGE to be a counterexample to NMIE, it must also be the case that Betty would not have merited blame for her intervention if she was right that the bridge had not frozen over. I presented and discussed SAFE BRIDGE to show that this is the case. While Angela does everything that Betty does in terms of her inquiry, the crucial difference is that, in SAFE BRIDGE, Angela comes to know that the bridge has not frozen and that taking the Risky Route is better for all. Since Angela knows that violating a coarse-grained, reasonable community norm favoring prudence would, in her case, be in everyone's best interests, there's little reason to think she merits blame in SAFE BRIDGE.

Since Betty seems to merit blame for her intervention even though she fulfills NMIE's conditions, FROZEN BRIDGE is a telling potential counterexample. For the purposes of this paper, it is not essential that we agree on a general account of all cases in which non-moral ignorance exculpates. Instead, what is important is to agree on one feature of FROZEN BRIDGE that is shared by certain cases in which moral ignorance also fails to exculpate. If we agree that these cases share this feature then, as will become clear in

[13]For further discussion on the fine-grainedness of rules, see Scanlon (1998: 198-206); for related discussion of the epistemic accessibility of the rules it is reasonable to expect others to follow, see Kumar (2015: 46). On internalism and rule-following, see Srinivasan (2015).

[14]Instead, what seems more likely to be true is that this is a norm that is meant as a rule of thumb, perhaps to be followed as a default. If this rule ought to be followed as a default, then this can also help explain why Betty merits criticism for violating it. In other words, the rule is one meant to be followed except under very special conditions, and FROZEN BRIDGE fails to satisfy those special conditions.

https://doi.org/10.1017/epi.2025.10056 Published online by Cambridge University Press

Section 3, we will be more open to my more general claim that, contrary to orthodoxy, there are striking parallels between the exculpatory power of both moral and non-moral ignorance.

The basic feature of FROZEN BRIDGE that is shared by cases in which non-moral ignorance fails to exculpate—cases that are discussed in-depth in Section 2—is that the agent violates a coarse-grained, reasonable community norm without knowing that doing so is best for everyone. When an agent violates such a community norm, she ought to take seriously the possibility that doing so is not in everyone's best interests. Betty has an action available to her that would not violate this community norm—taking the Prudent Route—so it seems perfectly fair for her to have to shoulder the moral consequences for her passengers' deaths. That is, if Betty violates this reasonable general norm without knowing that it's for the best, we seem to be on solid ground for blaming her for taking the Risky Route.[15]

To get a better handle on the basic structure of the argument, it'll be helpful to work through another pair of examples.[16]

> MEDICAL HURT: Emma comes across Fadi, who is seriously injured and waiting for an ambulance to arrive. If Fadi waits until the ambulance arrives to receive proper treatment, he will lose the ability to use one of his arms. If Fadi receives the proper treatment immediately, he will suffer so permanent injuries. If Fadi receives the improper treatment immediately, he will be paralyzed. Emma, who dropped out of medical school without a degree, and thus is not licensed to practice, knows all of this. After doing an incredible amount of research, Emma forms a false belief about the proper course of treatment for Fadi. Emma administers this treatment, resulting in Fadi becoming paralyzed.

In MEDICAL HURT, I think it's clear that Emma merits blame for administering improper medical treatment to Fadi. Roughly, the coarse-grained, reasonable community norm this seems to violate is something like, "Don't administer medical treatment to others without a medical license if professional help is on its way."[17] While all violations of norms are not blameworthy, when Emma violates this norm, she does not know that it will be in Fadi's best interests. Otherwise put, even though administering treatment has the potential for a huge moral payout—preventing permanent injury to Fadi—it was both more prudent for her to wait for the ambulance and she didn't know that she could do better herself. For this, it seems clear, Emma merits blame.

We can contrast MEDICAL HURT with the following:

> MEDICAL HELP: Gemma, similar to Emma, comes across Fadi who is seriously injured and waiting for an ambulance. Gemma also dropped out of medical school and lacks a license to practice. However, unlike Emma, Gemma's research results in her coming to know the proper course of treatment for Fadi.[18] Gemma administers this treatment, resulting in Fadi avoiding permanent injury altogether.

As discussed above with respect to Angela in SAFE BRIDGE, Gemma does not seem to merit blame in MEDICAL HELP. Since Gemma knew that, by administering the

[15]For discussion of related issues in the context of the use of defensive force, see Quong (2020: 39, 162).

[16]These are variations on a case discussed in Zimmerman (1997).

[17]Again, we might reject or reformulate this norm. Regardless of what norm we land on, however, a structurally similar case might be constructed.

[18]And, for those who may think it matters, she knows that she knows the proper course of treatment.

https://doi.org/10.1017/epi.2025.10056 Published online by Cambridge University Press

treatment, Fadi would emerge without permanent injury, she does not seem to merit moral criticism for administering the treatment. After all, if she simply waited for the ambulance, Fadi would have lost the use of one arm. Since Gemma knew that administering the proper treatment posed no risk to Fadi, we seem to lack solid grounds upon which to blame her for doing so. Otherwise put, Gemma knew that her course of action was best for everyone, so her violation of a coarse-grained, reasonable community norm does not seem blameworthy.

If Emma merits blame for her intervention in MEDICAL HURT, this puts further pressure on NMIE. This is because, like in FROZEN BRIDGE, it seems like Emma fulfils all of NMIE's conditions. First, Emma may have done everything we could reasonably expect of her to inquire into the proper treatment for Fadi. Second, the seeming truth of the proper treatment for Fadi may have played the appropriate role in Emma's decision to intervene. Lastly, if Emma was right about the proper treatment for Fadi, as we saw was the case with Gemma in MEDICAL HELP, she wouldn't have merited blame for intervening.

The feature of FROZEN BRIDGE that is shared by MEDICAL HURT, and that I return to below, is that Emma violated a coarse-grained, reasonable community norm without knowing that it was in everyone's best interests. In particular, in MEDICAL HURT, Emma ought to take seriously the possibility that, by administering treatment to Fadi, she might paralyze him. Since Emma has the option of being more prudent and following the reasonable general norm of waiting for an ambulance, it seems fair for her to shoulder blame for Fadi's catastrophic outcome.

The main takeaway from this discussion is that there are cases in which NMIE seems to fail. In particular, in cases in which agents violate coarse-grained, reasonable community norms without knowing that doing so is best for everyone, their non-moral ignorance fails to exculpate. This is a crucial finding since, as discussed below in Section 2, this feature is shared by cases in which moral ignorance fails to exculpate as well.

### *1.3. Potential objections*

Before exploring the parallels between non-moral and moral ignorance, it's worthwhile to pause and consider a couple of salient objections.

First, we might think that the account outlined above reduces blameworthiness to bad results. According to this objector, it's not fair to blame one person but not the other simply because of how things turned out. Otherwise put, if two agents perform the same action under the same circumstances, and one gets lucky good results while the other does not, then this is not enough to distinguish them in terms of their blameworthiness. Instead, says this objector, they ought to be equally blameworthy regardless of how things turn out. For the sake of argument, we can agree that our evaluations of others ought not reduce to whether their actions happen to turn out well. At the same time, the difference between Angela and Betty runs much deeper. In particular, it's not just the case that things turned out well for Angela, it's that Angela knew, beforehand, that her action would be best for everyone. If Angela did not know this, then perhaps she ought to be blamed just as much as Betty for taking unjustified risks with those in her care. However, since Angela did achieve knowledge, she did not take a "risk" with her passengers in the same was as Betty did, and it's not simply the case that, independent of anything about Angela, the universe smiled on her and things happened to go her way.

A second objection would be that an agent's blameworthiness ought to be indexed to her evidence. If we think that moral evaluation is subjective in this sense, *and we do not think that knowing a proposition is sufficient for that proposition to be part of one's evidence*, then the agents' evidence may be the same in both FROZEN BRIDGE and

https://doi.org/10.1017/epi.2025.10056 Published online by Cambridge University Press

SAFE BRIDGE, as well as MEDICAL HURT and MEDICAL HELP. If these agents' evidence is the same, and they perform the same actions based on this evidence, then it will be misguided, according to this objection, to blame one but not the other.

My first response to this objection is that, even if an agent's knowledge is not necessarily part of their evidence, it might still be the case that an agent who achieves knowledge of a proposition is in a stronger epistemic position than one with a false belief. If an agent who achieves knowledge is in a stronger epistemic position than one who does not, then it still seems to be the case that the two agent's evidence is not identical. My second response would be that, even if we are willing to abandon the position that knowledge is a uniquely strong epistemic state, we still will have some reason to blame Betty, but not Angela, for taking the Risky Route. Even from a position of epistemic parity, agents such as Betty ought to realize that, in taking the Risky Route, they are violating a coarse-grained reasonable norm for community behavior. With this fact in mind, she should carefully consider whether she is acting in a manner that will be best for all. Since Betty has an action available to her that does not violate these communal norms, it does not seem unfair to allow the moral consequences to fall on her shoulders if she places other's lives at risk. As discussed above, since Angela knew that her action was best for all—and that it would result in no deaths—it does not seem accurate to describe Angela as putting her passengers' lives at risk in the same way. In other words, since Betty knew she had less risky options open to her, it seems perfectly fair to blame her for violating community norms in a manner that resulted in the deaths of her passengers.

My third and last response to this objection—which is perhaps the most contentious—is that we should not expect that the justification that is available in the good case to be available in the bad case. In SAFE BRIDGE, Angela's action is justified by the fact that she knew that the bridge was not frozen and her action was in the best interests of all her passengers. In FROZEN BRIDGE, this justification is certainly not available to Betty. At best, Betty's action is justified by a sincere (but false) belief that she bridge is not frozen. In MEDICAL HELP, Gemma's action is justified by the fact that Gemma knows that her intervention will prevent Fadi from sustaining permanent injuries. In MEDICAL HURT, Emma's intervention is again, at best, justified by a false but sincere belief that she is helping as opposed to causing more harm. While there may be other justifications that we can offer in the bad cases, what will not be available is the justification that, by acting, the agent did what they knew was in everyone's best interests. If these agents' actions are justified in a distinct manner in the good and bad cases, then we will be unable, as this objector hopes, to draw a straight line from the agents' evidence to the justifications of their actions.

Even if we accept that there are counterexamples to NMIE, we may question their implications for moral ignorance. It is to these implications that I now turn.

## 2. When moral ignorance exculpates

To say the least, the claim that moral ignorance exculpates in the same manner as non-moral ignorance is contentious. As will become clear below, a central reason why we resist the exculpatory power of moral ignorance is that we do not want our ability to blame others for violating our rights to depend on their ability to figure out that doing so is wrong. To take a simple example, we want to preserve our right to blame those who intentionally cause us harm in order for the enjoyment it brings them, even if, for whatever reason, they fail to realize that such an action is morally abhorrent.

In this Section, I show how we can make space for the exculpatory power of moral ignorance without losing the ability to blame others when they violate coarse-grained,

https://doi.org/10.1017/epi.2025.10056 Published online by Cambridge University Press

reasonable community norms. To this end, I first discuss how the cases of moral ignorance that have been most discussed in the literature are ones that involve the violation of general community norms without knowledge that doing so is best for everyone. Since, according to my view, even non-moral ignorance doesn't exculpate in such cases, we can readily accept the parallel claim with respect to moral ignorance. To follow, I argue that, at least in some cases that do not involve the ignorant violation of general community norms, moral ignorance does, in fact, exculpate.

### *2.1. Moral ignorance and coarse-grained, reasonable community norms*

The terrain on which the battle over moral ignorance has been fought thus far is, I submit, misleading. This is because the cases that have been discussed are ones in which moral ignorance's failure to exculpate can be explained by the fact that agents violate coarse-grained, reasonable community norms without knowing that it's best for everyone. As argued above, when this happens in cases involving non-moral ignorance, exculpation does not occur. If we find that, in such cases, moral ignorance also fails to exculpate, this does not reveal some fundamental truth about moral ignorance's inability to exculpate; instead, it might just tell us something about ignorance's inability to exculpate when agents violate reasonable norms without the proper sort of justification.

Let's examine some of the cases that have been discussed in the literature. Perhaps the case that has received the closest scrutiny is that of a slaveholder in Biblical times.[19] Now, for the sake of argument we can assume such a slaveholder to be non-culpably ignorant of the wrongness of his action.[20] Even with this assumption, however, the fact that non-moral ignorance fails to exculpate in such a case can be explained by the fact that, in owning slaves, he violates a coarse-grained, reasonably community norm without knowing that it's best for all. The norm in this case might be something like "Respect other's autonomy" or even "Treat everyone as equals." In violating this general norm, the slaveholder ought to give very special consideration to the possibility that he is doing something pretty horrific. Otherwise put, the slaveholder ought to take very seriously the idea that his belief about the permissibility of owning slaves is false, just like Betty in FROZEN BRIDGE ought to take very seriously the idea that her belief that the bridge is frozen is false. A similar analysis can be applied to the case of the 1950s sexist who provides unequal resources and opportunities to his son and daughter. Even if we assume that his ignorance of the wrongness of sexism is non-culpable, it is still a case in which he violates a coarse-grained, reasonable community norm without knowing that its for the best. The general norm here could be something like "give equal opportunity to your children" or "support your children in developing to their full potential." Insofar as the sexist is violating norms in this manner, it seems fair to saddle him with the moral consequences if his justificatory beliefs turn out to be radically mistaken.

One objection to this analysis that is worth considering is that, if an agent's (for instance) racism is excused on account of moral ignorance, might his failure to grasp the relevant coarse-grained community norms also be excused? It is on this point that the account provided here draws a crucial distinction between course-grained and fine-grained norms—a distinction mirrored if not always made salient in the rest of the literature. Some types of what we might call 'fine-grained' moral ignorance—for instance ignorance of the equality of the sexes, the races, or of people with different nationalities—are often argued to provide a certain amount of exculpation. At the same

[19]Discussed most prominently by Rosen (2003, 2004, and 2008).

[20]Particularly instructive in this regard is Harman (2011) and Guerrero (2007); see also Fitzpatrick (2008) and Moody Adams (1994).

https://doi.org/10.1017/epi.2025.10056 Published online by Cambridge University Press

time, those who simply fail to recognize the common humanity of others—those who are ignorant of more 'coarse-grained' norms—are generally placed in a distinct category, perhaps akin to sociopathy or psychopathy. With this basic distinction in mind, it is just ignorance of the first kind of ignorance that is under consideration for being excused, while a complete lack of sympathy or empathy with others—coarse-grained ignorance—is not. In this way, according to the account provided, we can explore whether instances of fine-grained norm violations qualify for exculpation while insisting that some basic recognition of coarse-grained community norms is taken for granted.

I accept that, if moral ignorance does not exculpate as a general matter, this is a more informative explanation of why it doesn't exculpate when agents violate reasonable community norms without knowing that it's for the best. Ultimately, then, whether the reader should accept my explanation above will hinge on whether I can convince her that, in a range of cases in which such norm violations do not take place, moral ignorance does, in fact, exculpate. It is to this task that I now turn.

### *2.2. Moral ignorance with knowledge of what's best for all*

The claim I will defend in this section is that moral ignorance exculpates when the agent, although violating a coarse-grained, reasonable community norm, knows that doing so is best for all. Structurally, these cases will resemble SAFE BRIDGE, in which Angela violates such a norm while knowing that all of her passengers will be better off for her doing so.

Consider the following:

> FLOOR VOTE: The year is 1925. When it is up for debate, Senator Thomas casts the deciding vote against the Equal Rights Amendment (ERA), which would guarantee equality of rights for all persons regardless of sex. While Senator Thomas knows that all people of a certain standing deserve equal treatment, he doesn't believe that women possess that standing. If the senator did believe in the equality of the sexes, he would've voted for the ERA. Senator Thomas knows that, despite his opposition, the ERA will pass the following year. Lastly, he knows that women will be better off if the ERA passes the following year because the revised bill will be even stronger.

In FLOOR VOTE, although Senator Thomas violates a coarse-grained, reasonable community norm—perhaps one mandating equal treatment for all—he knows that doing so is ultimately best for everyone. More specifically, although senator Thomas is ignorant of the moral equality of men and women, he still acts in a manner that he knows will be most beneficial to women. As such, he ought not be blamed for voting against the ERA. It certainly is the case that Senator Thomas is on the wrong side of history, and he may completely fail to understand what his moral obligations are in such a situation. At the same time, insofar as the Senator knows that everyone (including women) will be better off for his opposition, he doesn't seem to merit blame for what he does.[21]

Two objections are worth addressing at this point. First, one might object that certainly the Senator merits some moral criticism in FLOOR VOTE. More specifically, the Senator seems to merit criticism for failing to believe that women are equal to men. Insofar as we want to criticize individuals for their attitudes *independently* of the manifestation of those attitudes in action, I need not disagree. At the same time, I am

[21]For an argument that blame is such cases ought to be graded, see Sliwa (2010).

https://doi.org/10.1017/epi.2025.10056 Published online by Cambridge University Press

much more concerned with the blame and moral criticism that agents merit for their actions. At a minimum, we can intelligibly ask whether agents are excused for the manner in which they treat each other independent of whether, in some sense, their underlying attitudes and prejudices have serious shortcomings.

Another objection that we might have is that Senator Thomas would only be excused for his actions in FLOOR VOTE if his knowledge played the appropriate role in the production of his action. In particular, we might insist, Senator Thomas would only be excused if it was the case that, if he did not know that voting against the ERA would be in everyone's interests, then he would not have done so. While I have sympathy with this reaction, for the purposes of our examination, it is more important that we require of the Senator that, if he knew that women had equal status to men, he would've supported the ERA. That is, according to MIE, to even be a candidate for being exculpated by moral ignorance, it must be the case that one's moral ignorance played the appropriate role in the production of one's action. Since, in order to be relevant to our evaluation of MIE, the Senator's vote must turn on whether he believes in women's equality, it cannot also turn on his knowledge that the vote is in everyone's best interests.[22]

In FLOOR VOTE, I argued that the Senator's moral ignorance ought to exculpate him for his failure to support the ERA. Even though he violates a coarse-grained, reasonable community norm, it is also the case that he knows that doing so is in everyone's best interests. This is a similar analysis that we applied to cases in which non-moral ignorance exculpates. For instance, in both SAFE BRIDGE and MEDICAL HELP, the relevant agents violated coarse-grained, reasonable community norms with the knowledge that doing so was for the best. While all agents ought to take the prospect of violating reasonable community norms quite seriously in their deliberation, they can rest assured that, even if they violate such norms, they will be in the clear if they know their actions are best for everyone. At the same time, as illustrated in FROZEN BRIDGE and MEDICAL HURT, these agents ought to remember that, if they lack this knowledge, and their actions end up harming others, then the violation of general norms may render them blameworthy.

To help solidify ideas, one more example is worth considering:

> HARSH MENTOR: Seventy-five year old Professor Penny is a detached, strict, and demanding advisor to Student Sam. In Professor Penny's view, harsh feedback helps Student Sam grow and develop, and she also knows that Sam does not take her sharp comments personally. While some call Professor Penny old fashioned, she thinks that her techniques are the most effective and just. In addition, she knows that her treatment of Sam is in his best interests not only academically but also personally, as a harsh coach is the exact sort of coach that Sam loves and thrives under. Contrary to Professor Penny's beliefs—and as she is later informed by her university administrators—her brusque and unforgiving manner of treating Sam is morally impermissible.

In HARSH MENTOR, like in FLOOR VOTE, although Professor Penny violates a reasonable community norm—perhaps one mandating that professors treat students with a minimal amount of care, respect, and courtesy—her ignorance ought to exculpate

[22]In addition, it is worth noting that I think the evaluation of counterfactuals (here and elsewhere) is far from straightforward. In particular, in the counterfactual case in which voting against the ERA was not in everyone else's interests, what exactly would have made that the case? Perhaps more importantly, why is it relevant to our evaluation of the Senator's actual actions what would have been the case in nearby, but non-actual, possible worlds? For further discussion on this point, see Sliwa (2016) and Hills (2009).

https://doi.org/10.1017/epi.2025.10056 Published online by Cambridge University Press

her. This is because, although she is violating this norm, and perhaps ought to know better, she also knows that she is acting in Sam's best interests. Since she knows she is doing what's best for Sam, we ought not blame her for failing to adhere to relevant community norms.

Like FLOOR VOTE, cases like HARSH MENTOR support my claim that moral ignorance can exculpate in cases in which agents, though violating coarse-grained, reasonable community norms, know that their actions are best for all. Before closing, however, one last objection is worth addressing. In particular, in a case such as HARSH MENTOR, we might wonder whether it is the ignorance that is exculpating or the fact that Professor Penny is acting in Sam's best interests. Otherwise put, we might wonder whether agents ever merit blame for acting in a manner that is everyone's best interests. If not, then we have an alternate account of why such agents are not blameworthy for their actions. While this objection may be tempting for some, I think it will fall short for those who do not think that the moral can be reduced to the promotion of individual interests. For instance, in FLOOR VOTE, I think there is a strong case to be made that the morally right course of action is to vote in favor of equality, even if, from a strategic perspective, everyone might be better off with a different vote. Along similar lines, in HARSH MENTOR, I think there is a strong case to be made that we always ought to treat our students with respect, even if they might benefit from some 'tough love.' These particular verdicts notwithstanding, however, the moral general point can surely be made that, unless we are enamored with consequentialism, our moral obligations to others might not always align with acting in a manner that promotes individual interest. In whatever range of cases this may be, agents might be exculpated in virtue of their ignorance and this exculpation will not stem from the fact that they act in a manner that's best for all.

## 3. Conclusion

In this paper, I argued that, in a range of important cases, non-moral ignorance does not exculpate. More specifically, I showed that non-moral ignorance does not exculpate when an agent violates a coarse-grained, reasonable community norm without knowing that doing so is in everyone's interests. To follow, I argued that moral ignorance, like non-moral ignorance, does exculpate when such norms are not violated in this manner. With these two claims in place, we arrive at a coherent picture of one way in which either moral or non-moral ignorance can exculpate. According to this picture, the same basic factors are at play in determining whether non-moral and moral ignorance exculpates. First, we determine whether a coarse-grained, reasonable community norm was violated, and second, we determine whether the agent knew that such a violation was in everyone's best interests. If the answer to both questions is 'yes', we have a prime candidate for blameless ignorance.

## References

**Anderson, C. and Hawthorne, J.** (2019). '*Pragmatic Encroachment and Closure,' in Pragmatic Encroachment in Epistemology. Routledge Studies in Epistemology*. New York: Taylor & Francis.

**Aristotle** (2009). Nicomachean Ethics, translated by David Ross, revised by Lesley Brown. Oxford: Oxford university Press, Oxford World's Classic.

**Brown, J.** (2008). 'Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning.' *Nous* **42**(2), 167–89.

**Fitzpatrick, W. J.** (2008). 'Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge'. *Ethics* **118**(4), 589–613.

https://doi.org/10.1017/epi.2025.10056 Published online by Cambridge University Press

**Guerrero, A.** (2007). 'Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution.' *Philosophical Studies* **136**, 59–97.
**Harman, E.** (2011). 'Does Moral Ignorance Exculpate?' *Ratio* **24**(3), 443–68.
**Hills, A.** (2009). 'Moral Testimony and Moral Epistemology.' *Ethics* **120**(1), 94–127.
**Kumar, R.** (2015). 'Risking and Wronging.' *Philosophy and Public Affairs* **23**(1), 27–51.
**Moody-Adams, M.** (1994). 'Culture, Responsibility, and Affected Ignorance.' *Ethics* **104**, 291–309.
**Quong, J.** (2020). *The Morality of Defensive Force*. Oxford: Oxford University Press.
**Rosen, G.** (2003). 'Culpability and Ignorance.' *Proceedings of the Aristotelian Society* **103**(1), 61–84.
**Rosen, G.** (2004). 'Skepticism about Moral Responsibility.' *Philosophical Perspectives* **18**, 295–313.
**Rosen, G.** (2008). 'Kleinbart the Obvious and Other Tales of Ignorance and Responsibility.' *Journal of Philosophy* **105**(10), 591–610.
**Scanlon, T.M.** (1998). *What We Owe Each Other*. Cambridge, Mass: Harvard University Press.
**Sliwa, P.** (2016). 'Moral Worth and Moral Knowledge.' *Philosophy and Phenomenological Research* **93**(2), 393–418.
**Smith, H. M.** (1983). 'Culpable Ignorance.' *The Philosophical Review* **92**(4), 543–71.
**Sridharan, V.** (2020). 'Two Faces of Rationality.' *Synthese* **198**, 11103–24.
**Sridharan, V**, (2024). 'On the Right to be Practically Sure,' Inquiry, https://doi.org/10.1080/0020174X.2025.2470870
**Srinivasan, A.** (2015). 'Normativity without Cartesian Privilege.' *Philosophical Issues* **25**(1), 273–99.
**Williamson, T.** (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
**Zimmerman, M. J.** (1997). 'Moral Responsibility and Ignorance.' *Ethics* **107**, 410–26.

**Vishnu Sridharan** is an Assistant Professor of Philosophy at the University of Colorado, Boulder. He recently completed a Postdoctoral Fellowship in Law and Philosophy at UCLA. He received his PhD from the University of Southern California and his J.D. from Stanford Law School. He can be reached at Email: vishnu.sridharan@colorado.edu.

**Cite this article:** Sridharan V. (2025). "When it is (and is not) Blameworthy to Break the Rules." *Episteme* 1–14. https://doi.org/10.1017/epi.2025.10056