

Recklessness and Uncertainty: Jackson Cases and Merely Apparent Asymmetry

(Forthcoming in *Journal of Moral Philosophy*)

Claire Field

Abstract:

Is normative uncertainty like factual uncertainty? Should it have the same effects on our actions? Some have thought not. Those who defend an asymmetry between normative and factual uncertainty typically do so as part of the claim that our moral beliefs in general are irrelevant to both the moral value and the moral worth of our actions (Weatherson 2014; Harman 2015). Here I use the consideration of Jackson cases to challenge this view, arguing that we can explain away the apparent asymmetries between normative and factual uncertainty by considering the particular features of the cases in greater detail. Such consideration shows that, in fact, normative and factual uncertainty are equally relevant to moral assessment.

Keywords: Jackson cases; normative uncertainty; moral recklessness; moral appraisal

1. Introduction

Can normative uncertainty affect what morality requires of us? If it is anything like factual uncertainty, it can. However, some have thought that there are significant differences between the two. Those who defend these differences typically do so as part of the claim that our moral beliefs are in general irrelevant to both the moral value and the moral worth of our actions (Weatherson 2014; Harman 2015). Jackson cases offer us a way to evaluate the truth of this claim about the asymmetry between normative and factual uncertainty. Jackson cases are situations in which the agent is uncertain about the moral values of the possible outcomes in the following way:

Jackson Case:

	Beliefs	Facts
A	Either +100 or -100	+ 100
B	0	0
C	Either +100 or -100	-100

There is widespread agreement that someone appropriately concerned for moral value would choose option B¹. This is because even though more moral value could be achieved by choosing A, to do this given what the agent currently believes about A would be reckless, and people appropriately concerned with moral value do not gamble with it in this way. There is debate to be had over whether this implies that option B is really the morally right thing to do, as Zimmerman (1997, 2008, 2014) argues, or whether really option A is the morally right thing to do, but there is some other consideration that explains why it would not be morally conscientious to choose it in this case. For example, according to Bykvist (2014), option A is the morally right thing to do, but because the morally conscientious person has a preference for acting so as to promote moral value and avoid moral disvalue it would be rational for her to choose option B, because this maximises expected moral value. Similarly, Graham (2010) claims that the morally conscientious person prudentially ought to choose option B, even though morally she ought to choose option A; and that furthermore it

¹ See for example, (Zimmerman 1997; Zimmerman 2008; J. 2014; Bykvist 2014; Graham 2010; Sepielli 2009; Moller 2011).

would be blameworthy to choose option A, even though this is the morally right option. Despite their differences, what all these views can agree on is that given the agent's beliefs, it would be reckless for her to choose anything other than option B, and this recklessness deserves a negative moral assessment.

It might be thought that if choosing anything other than option B is reprehensibly reckless, then this is true regardless of what it is that the agent is uncertain about. Harman (2015) and Weatherson (2014) disagree. They argue that choosing either option A or C does not deserve negative moral assessment if the uncertainty involved is normative - that is, if it is uncertainty about the moral deontic status of the possible outcomes - whether they are morally permissible, forbidden, or required. They argue that this is because our beliefs about the moral deontic status of possible outcomes are irrelevant to moral evaluation - unlike other beliefs, they do not affect what we morally ought to do, whether or not we are blameworthy, or the moral worth of our actions. This paper challenges this supposed exceptionality of normative uncertainty, arguing that once we factor in the way that what one ought to do depends on the exact degree to which we are uncertain, the stakes involved in the case, and the quality of the agent's evidence, we can explain away any apparent distinctiveness of normative uncertainty. This provides some motivation to think that our beliefs about what is morally permissible, forbidden, or required are relevant to moral evaluation in just the same way that other beliefs are.

2. Jackson Cases

The original Jackson case involves factual uncertainty:

Drugs. A doctor has a patient with a minor but not trivial skin complaint. The doctor has three drugs to choose from, and careful consideration of the medical evidence available to her has led her to the following opinions. One of drugs A or C will completely cure the skin condition, the other will kill the patient. Crucially, she has no way of telling which of A and C will cure the patient, and which will kill him. However, she knows that drug B is very likely to relieve the condition but will not completely cure it (see Jackson, 1991: 462).

It seems clear that the doctor should prescribe drug B, and this is despite the fact that she knows it is not the best option. For her to prescribe either of the other two drugs would be reckless. Prescribing either of drugs A or C would be reckless because her available information does not allow her to discriminate between them; for all she knows she could be prescribing poison, and curing a minor skin condition is not worth the risk of death. We can represent the case as follows:

Drugs	Beliefs	Facts
A	Death or Cure	Cure
B	Partial Cure	Partial Cure
C	Death or Cure	Death

Jackson cases meet the following criteria:

- The agent morally ought to choose an option that is sub-optimal.
- The agent knows that the option she ought to choose is sub-optimal.

- It would be unacceptably reckless for the agent to choose any other option.

In *Drugs*, the doctor is uncertain about some of the non-moral facts of the situation. Analogous cases involving evaluative uncertainty and normative uncertainty are also possible. An agent who is evaluatively uncertain is uncertain about the what the moral values of her options are², she is uncertain whether performing the action would be morally good or morally bad. An agent who is normatively uncertain is uncertain about the deontic status of her options³, she is uncertain whether performing the action would be morally permissible, impermissible, or required. It is possible to be evaluatively uncertain without being normatively uncertain – one might be uncertain about the moral value of an outcome while being completely sure that it would be morally impermissible to perform the action, because doing so would be to recklessly risk doing something very bad. In other cases, agents will be both evaluatively and normatively uncertain, since if one is uncertain about the moral values of a case, then this can produce uncertainty about whether or not the action is morally permissible to perform. The focus of this discussion is on normative uncertainty. In the cases under consideration, the agents are normatively uncertain because they are uncertain about which is the correct moral theory, and the competing moral theories disagree about the moral values of each option.

² See Zimmerman (2008, 2014) for some discussion.

³ See Graham (2010) for some discussion.

Like factual Jackson cases, normative Jackson cases should involve a choice between three options that are comparable to those of the original case: one should be a solution that will resolve the situation perfectly, producing the best outcome; another should be a catastrophe, and the third should be a partial resolution that does not resolve the situation completely. As in the original case, the agent should have no way of distinguishing between the perfect solution and the moral catastrophe. Here is an example:

Voting. Bill is trying to decide which political party to vote for. He could vote for the Equality party, the Freedom party, or the Other party. Everyone he knows strongly supports the Equality party, but he is not entirely convinced by their emphasis on equality. He is also not entirely convinced by the policies of the Freedom party, but he finds some of them very persuasive. He has studied some political philosophy and finds himself pulled in two opposing directions. On the one hand he suspects the ideals of libertarian freedom upheld by the Freedom party are valuable, and knows that these are incompatible with the ideals of equality upheld by the Equality party. On the other hand, he is persuaded by the idea that equality can sometimes promote more happiness for the underprivileged, and wonders if libertarianism's emphasis on freedom is misguided. He is entirely unmoved by the policies of the Other party, who do little to promote either freedom or equality, but he knows that were they to be elected they would at least do no great harm and are likely to make some small improvements in the efficiency of public services. He must vote tomorrow, and has no more time to deliberate.

Assume that Bill lives in a political system where his vote will have a direct effect on which party will be elected, such that which way he votes is a genuinely moral question. Each of the extremist parties would implement policies that promote one and only one of the political ideals – libertarian freedom or equality in distribution. The Other party would promote other, comparatively negligible, goods such as the

efficient running of public services. The political situation that would result from the election of either of the Freedom party or the Equality party would be either optimal or a disaster, depending on which moral theory turns out to be true, while the election of the Other party would have a neutral outcome. We can represent the case as follows:

Voting	Beliefs	Facts ⁴
Equality Party	Either very morally good or very morally bad.	Very morally good.
Other Party	Neither morally good nor morally bad.	Neither morally good nor morally bad.
Freedom Party	Either very morally good or very morally bad.	Very morally bad.

Whether voting for the Equality Party would have a positive or negative value moral depends on the moral facts of the case – on whether it is equality or freedom that is more morally valuable in this case. These moral facts affect what Bill ought to do when he goes to the voting booth, and it is these moral facts that he is uncertain about.

It is compatible with this being a genuine Jackson case that both freedom and equality are indeed political goods, and promoting them would be morally good in some situations; but in this situation the outcome of promoting one would be a catastrophe,

⁴ The reader is free to switch the values of the actual outcome to suit her own political inclinations.

while the outcome of promoting the other would be best. Alternatively, it may also be the case that only one of freedom or equality is a genuine good, or that both are goods but one is significantly more important than the other. The important point is that Bill cannot tell from his current information whether it is freedom or equality that is required in this situation, but he knows that one would be the morally best option, one would be the morally worst option, and there is a third morally safe but sub-optimal option available to him. If this is really a Jackson case then Bill's uncertainty is relevant to what he ought to do, and he ought to choose the least risky option given his current evidence, and vote for the Other party⁵. Here is another case:

Amputation. Tanishka is a doctor. She has a patient who is presenting symptoms of Body Integrity Identity Disorder. Patients with this disorder typically claim that a healthy and normally functioning limb is not theirs, that it feels 'strange', and that they want it to be removed. Post-amputation, patients typically report feelings of relief. Furthermore, Tanishka knows that patients with BIID who do not receive amputation often attempt to amputate limbs themselves, usually botching the job and risking their lives. Tanishka's patient has asked Tanishka to remove his otherwise healthy left leg. Like most doctors who have taken courses in medical ethics, Tanishka believes both that doctors ought to respect the autonomy of their patients, and that doctors ought to avoid causing their patients harm. She believes that in order to respect her patient's autonomy she should amputate the leg, but she also believes that the amputation would constitute causing her patient harm. She believes that if autonomy is of more moral importance then morally she ought to amputate, but if avoiding harm is of more moral importance then morally she ought not amputate, and she ought to take steps to ensure that the patient is prevented from performing his own amputation and botching it. She has a third option available; she could refer her patient for psychiatric therapy. She knows that such therapy

⁵ See (Guerrero 2007; Lockhart 2000; Sepielli 2009; Moller 2011) for other defences of mitigating risk in situations of normative uncertainty.

would be unlikely to completely cure the disorder, but it is likely to relieve some of the symptoms of the condition.

Tanishka's uncertainty is uncertainty about the relative importance, in this situation, of respecting autonomy and avoiding causing the patient harm. She knows that both are important, and she knows that they conflict in this situation. What she does not know is which she ought to be prioritise. If she amputates, she either does a good thing by respecting the patient's autonomy, or she does a very bad thing in causing him harm. If she does not amputate, and keeps the patient under surveillance then she either does the right thing in preventing harm, or she does a very bad thing in unjustifiably constraining the patient's autonomy in two ways – by not respecting his wishes to have his leg amputated, and then by infringing his privacy and liberty by keeping him under surveillance. When Tanishka deliberates she is trying to decide on the relative moral importance, in this situation, of autonomy and harm avoidance. Like the agents in the Voting and Drugs, she knows that she has a safe option available. Referring the patient for psychiatric therapy carries no risk of either of the worst options, but also does little to help the patient's condition. We can represent the case as follows:

Amputation	Beliefs	Facts ⁶
Amputate	Either very morally bad or very morally good.	Very morally good.
Psychiatric therapy	Neither morally good nor morally bad.	Neither morally good nor morally bad.
No amputation, keep patient under surveillance.	Either very morally bad or very morally good.	Very morally bad.

Those who deny the possibility of Jackson cases deny that the agents in Voting and Amputation ought to take the low risk option. Harman (2015) and Weatherson (2014) take the view that our moral beliefs are irrelevant to moral conscientiousness in a way that factual uncertainty is not, and that therefore there are significant asymmetries between how normative and factual uncertainty affects whether or not it would be appropriate for someone concerned with morality to take the safe option. They take the view that it is not morally conscientious to risk great wrongs by acting under factual uncertainty, but it is also not morally conscientious to respond to normative uncertainty by choosing the low risk option⁷. This position is suggested by some

⁶ The reader is free to switch the values of the actual outcome to suit her own moral views if she wishes.

⁷ It is worth noting that Harman and Weatherson also share Zimmerman's reading of the original factual Jackson case, and take the view that the doctor morally ought to prescribe drug B, because it is wrong to risk the patient's life for the sake of a minor skin condition. I use the term 'morally conscientious' here so as to avoid taking a stand on what the agent morally ought to do. My quarrel is with the purported asymmetry between factual and normative uncertainty, rather than with any particular claim about what the agent morally ought to do.

apparent asymmetries between cases of normative and factual uncertainty. The following sections argue that these asymmetries can be explained away. Specifically, once we factor in the way that what it is appropriate for the morally conscientious person to do depends on the exact degree to which she is uncertain, the stakes involved in the case, and what her evidence supports, we can see that moral recklessness is to be avoided just as factual recklessness is, and there should be no asymmetry in our responses to Jackson cases involving normative and factual uncertainty.

3. The Case for Asymmetry

There are three apparent asymmetries between moral and factual uncertainty that the case for asymmetry rests on. If they are genuine, then there is reason to think that Harman and Weatherson are right that it would not be reprehensively reckless to choose the low risk option in normative Jackson cases, as it would in factual Jackson cases. In this section I present the three putative asymmetries, and in the following section I show how they can be seen to be merely apparent once we consider their specific features in more detail.

The first putative asymmetry is that there are cases in which it can seem morally permissible for an agent to perform actions that she is normatively uncertain about, while it is not permissible for her to perform actions that she is factually uncertain about. For example:

*Dinner*⁸. Martha is deciding whether to have steak or tofu for dinner. She prefers steak, but knows there are ethical questions around meat-eating. She has studied the relevant biological and philosophical literature, and concluded that it is not wrong to eat steak. But she is not completely certain of this; as with any other philosophical conclusion, she has doubts. As a matter of fact, Martha is right in the sense that a fully informed person in her position would know that in this instance it is permissible for her to order steak for dinner, but Martha cannot be certain of this.

We might represent the Martha’s options as follows:

Dinner	Eating meat is wrong		Eating meat is not wrong	
	Moral Value	Hedonistic Value	Moral Value	Hedonistic Value
Steak	-10	+20	0	+20
Tofu	0	+5	0	+5

Weatherson’s claim is that it is not wrong for Martha to have steak for dinner, despite the fact that she is completely certain that she has an option that is morally permissible – tofu⁹. If Martha chooses steak, then she risks incurring moral disvalue to gain the small pleasure of eating steak. Weatherson claims that intuitively, it would not be wrong for Martha to choose steak and this shows that moral recklessness cannot be wrong. If this is correct, then this would be one respect in which normative uncertainty differs from factual uncertainty. It would show that it is morally permissible to make reckless choices when the uncertainty is normative, but not when it is factual – as in

⁸ This example is borrowed from Weatherson (2014: 2).

⁹ See also Harman's (2016) view that eating meat is a suberogatory but not morally prohibited action, a ‘morally permissible moral mistake’.

the Drugs case. In §4.1 I argue that this is not the conclusion we should draw from consideration of this case, because there are alternative explanations for the intuition that Martha is permitted to choose steak.

The second apparent asymmetry is that there seem to be cases in which agents who act against their moral beliefs are praiseworthy, but if moral recklessness was a vice, then these agents would be the prime examples of morally reckless agents and we would expect them to be blameworthy rather than praiseworthy. Arpaly gives the example of the inversely akratic agent, Huck Finn:

Huck Finn. Huck Finn befriends Jim, a slave, and helps him escape from slavery. While Huck and Jim are together on a raft used in the escape, Huck is plagued by what he calls “conscience.” He believes, as everyone in his society “knows,” that helping a slave escape amounts to stealing, and stealing is wrong. He also believes that one should be helpful and loyal to one's friends, but loyalty to friends is outweighed by some things, such as property rights, and does Miss Watson, Jim's owner, not have property rights? Hoping against hope to find some excuse not to turn Jim in, Huck deliberates. Like many children (and adults), Huck is not very good at abstract deliberation, and it never occurs to him to doubt what his society considers common sense. Thus, he fails to find a loophole. “What has poor Miss Watson done to me,” he berates himself, “that I can see her [slave] go away and say nothing at all?” Having thus deliberated, Huck resolves to turn Jim in, because it is “the right thing.” But along comes a perfect opportunity for him to turn Jim in, and he finds himself psychologically unable to do it. He accuses himself of being a weak-willed boy, who has not “the spunk of a rabbit” and cannot bring himself to do the right thing, and eventually shrugs and decides to remain a bad boy (Arpaly 2002: 75).

We might represent Huck's options as follows:

Huck Finn	Freeing Jim is right	Freeing Jim is wrong
Turn Jim in	Very morally bad	Very morally good
Free Jim	Very morally good	Very morally bad

Intuitively, Huck's action is praiseworthy, and this is despite the fact that he does it in the belief that it is wrong. He pays attention to the right kinds of things, Jim's humanity and welfare, and this means that he cannot bring himself to turn Jim into the authorities, even though he mistakenly believes that turning him in would be the right thing to do. This seems to suggest that moral recklessness is not really a vice – Huck seems even more morally reckless than someone who fails to choose the low risk option when they are uncertain – he actually believes what he does is wrong, and does it anyway. If moral recklessness was a vice, then it seems we should blame him rather than praise him, and we should think he ought not free Jim. This apparent asymmetry is dealt with in §4.2.

The third apparent asymmetry is that the normative and factual versions of taking the low risk option do not seem to have equal moral value when considered as long term decision making policies. While it can seem like a good policy for agents to take the low risk option in all cases of factual uncertainty, the same is not obviously true for normative uncertainty. Agents who consistently choose the risk-minimising option can seem morally cowardly, and moral cowardice can seem just as much of a moral

vice as moral recklessness, if not more so¹⁰, and it seems that such cowardice would be ubiquitous given that most agents are almost never certain about their moral views. Taking the view that appropriate concern for morality requires choosing the option that mitigates risk whenever one is less than fully certain, and when this option is typically morally sub-optimal, means that there will be almost no cases in which agents with correct moral beliefs are permitted to do the right thing. Employed over time, this position advocates a state of moral inertia. If acting on less than certain moral beliefs were always wrong, then it can seem that we would be forced to view agents who campaign for moral change, against the dominant views of their communities, as morally reckless and so not morally conscientious. On the assumption that the opinions of supposed moral authorities count as evidence for and against moral views, agents such as Rosa Parks, a civil rights campaigner in 1950s America would count as examples of moral recklessness, and the more ordinary people who followed her perhaps even more so. In the face of so much disagreement, it can seem unlikely that they could have been sure that the new views were correct. So, in so far as acting on uncertain beliefs is reckless, they were not acting in a way that exhibited appropriate concern for morality. Cautiousness seems much less admirable when it is moral cautiousness rather than factual cautiousness – we typically evaluate advocates for moral change such as Rosa Parks as having acted in a morally admirable way, we see them as admirable examples of moral bravery, rather than reprehensibly reckless.

¹⁰ See Weatherson (forthcoming) for an articulation of this thought.

If acting under normative uncertainty was reckless in the same way that factual uncertainty was, then there would be no moral heroes. Agents who act rightly but against the majority opinion of their communities would always be acting recklessly – either they would be less than certain that they were acting rightly, because community disagreement would reduce their credence in their own rightness; or, if they did have credences high enough to justify action, then they would be doing so while recklessly disregarding evidence from community disagreement. So, if moral recklessness is to be thought of in the same way as factual uncertainty, then it seems we ought to prefer that these agents had pursued more conservative courses of action, and this seems unpalatable. However, §4.3 explains why this is not a genuine asymmetry.

4. Alternative Explanations

Defenders of the asymmetry between normative and factual uncertainty argue that what explains these asymmetries is that our beliefs about what is morally right have no effect either on what we morally ought to do, or on whether we are to be praised or blamed; they are irrelevant. This means that whether or not we are uncertain about moral claims has no effect on what we should do, and so moral recklessness is not a vice in the same way that factual recklessness is. I argue that this is the wrong conclusion to draw from the three apparent asymmetries outlined in the previous section, and suggest that closer consideration of the agent's epistemic situation and the stakes involved shows that these asymmetries are merely apparent. This insight

shows that moral and factual recklessness, are in fact symmetrical, and so we have no reason to think that moral recklessness is not a vice.

4.1 Dinner

Dinner was supposed to show that moral recklessness is not wrong, even on the assumption that factual recklessness is. This asymmetry was based on the intuition that it is permissible for Martha to eat steak even though in doing so she risks some moral disvalue. However, it is not clear that the case does show this, because there is more than one way to explain this intuition and the putative non-viciousness of moral recklessness is only one such way. This section outlines two alternative explanations of the case: one based the observation that morality does not necessarily require us to always maximise moral value at the expense of all other value, and another based on the difficulty in distinguishing when our intuitions are tracking blameworthiness and when they are tracking moral permissibility.

First, we might agree that Martha is permitted to choose the steak, but follow Harman (2016) in viewing this as morally suberogatory but permissible. For example, we might think although it would not be the morally best thing for her to do, she is permitted to risk a small amount of moral disvalue for the pleasure she would gain from eating the steak. This would mean that Weatherson is right about the moral permissibility of Martha's choosing steak, but wrong to take this to imply that moral recklessness is never wrong. Instead, the right explanation is that morality does not always require us to maximise moral value at the expense of all other kinds of value, and this is a case

in which morality permits us to prioritise pleasure over the risk of moral disvalue. There is no asymmetry with factual uncertainty here – various examples from daily life suggest that it is at least sometimes permissible to risk some moral disvalue for our own pleasure¹¹. For example:

Driving. Amber is deciding whether to visit the art gallery or stay at home. To get to the art gallery she must drive. She would prefer to visit the art gallery, but she cannot be certain that she will not kill someone on the way. She knows that there is a small possibility that she will harm someone while driving.

Driving	Someone is harmed		No one is harmed	
	Moral Value	Hedonistic Value	Moral Value	Hedonistic Value
Drive to art gallery	-1000	+50	0	+50
Stay at home	0	-50	0	-50

The case is structurally similar to Dinner, except that it involves factual rather than moral uncertainty. By driving to the art gallery, Amber is performing an action that she cannot be completely certain will avoid harm, because driving accidents are always possible. Nevertheless, risks of this kind are usually thought of as morally permissible in the factual case, even if not the morally best thing one could do. This is true even though the harm risked is even greater than in the Dinner case. This does not suggest an asymmetry between factual and normative uncertainty, but rather that

¹¹ See also Guerrero (2007) on this point.

there are simply some risks that it morally permissible to take, and this is true regardless of whether the uncertainty is factual or normative. If it is permissible to risk the possibility of causing a traffic accident for the small pleasure of going to an art gallery, then it should not be surprising that the even smaller risk of wrongdoing in Dinner is also morally permissible. However, if there really was nothing of any value to be gained from driving, or from eating the steak, then it is not so clear that we would have the intuition that these actions were permissible¹². Again, this is a respect in which factual and normative uncertainty are symmetrical.

Second, an alternative explanation of the intuition is that it is not tracking moral permissibility at all, but rather blamelessness. Weatherston interprets the intuition as one that indicates the moral permissibility of choosing steak, but it is not so clear that our intuitions are sufficiently fine-grained for us to reliably distinguish intuitions of moral permissibility from intuitions of other moral evaluations, such as blamelessness or excusibility¹³. An alternative explanation of the intuition is that, in fact, Martha is not permitted to choose the steak – she really ought to choose tofu – but we are unwilling to blame her of failing to do so. We might think that it is never permissible to prioritise one's own pleasure over any risk moral of disvalue, moral value is just too important to risk in this way. This would mean that neither Martha is not permitted to choose the steak, and nor is Amber permitted to drive to the art gallery,

¹² Guerrero (2007) also makes this point.

¹³ Various epistemologists take this to be a common phenomenon in our thinking about justification in the epistemic case (see (Littlejohn, forthcoming; Sutton 2007; Williamson 2017)).

but we might nevertheless be unwilling to blame her, and this unwillingness would explain the intuition that Weatherson draws on.

There are at least two features of the case that could explain our unwillingness to blame Martha, despite the fact that she does something that is strictly impermissible. One such feature is that we might think blame is a serious sanction that is appropriate only for relatively serious transgressions, in which more is at stake than just one steak at one dinner. Although Martha does make a reckless choice, it is only very slightly reckless. So slight, in fact, that describing the case as 'reckless' sounds a little odd; we typically reserve the term for more serious cases where greater risks are taken. Even on the assumption that eating meat is wrong, choosing steak once at one meal is not a great moral catastrophe, and certainly not a catastrophe comparable to death in the original Jackson case. As such, risking this minor wrong is perhaps not something deserving of blame, even if the agent ought not to have done so. It is plausible that our intuitions might change were we to increase the stakes:

Veganism. Martha is trying to decide whether to become a vegan or not. She enjoys eating meat and dairy products, so would prefer not to. She has studied the relevant philosophical literature carefully and concluded that it would not be morally wrong for her to choose to continue consuming meat and dairy products. However, she is not completely certain of this; as with any other philosophical conclusion, she has doubts. She also knows that if it turns out that she is wrong, the impact that a lifetime of consuming meat and dairy products would have would constitute a serious moral wrong. As a matter of fact, Martha is right in the sense that a fully informed person in her position would know that it is permissible for her to not become a vegan, but Martha cannot be certain of this.

The consequences of a lifetime of meat eating are much greater than the consequences of one steak at one dinner, and this is a case in which it is much less clear that Martha is blameless for choosing a lifetime of meat-eating, given her uncertainty.

Another reason we might be hesitant to evaluate Martha as blameworthy is that there are features of her epistemic situation that make her choice only very slightly reckless, and much less reckless than cases of recklessness that clearly would deserve blame, such as choosing anything but the partial cure in *Drugs*. Firstly, Martha is much more confident than a typical reckless agent. In *Drugs*, the doctor had ‘no way of telling’ which of drugs A and B was which, and this low confidence contributes to the sense that it would be very reckless for her to choose either of drugs A or B. However, Martha in *Dinner* is comparatively confident that it is permissible to eat steak – she is ‘not completely certain’, making her choice to eat steak much less reckless than if the doctor were to choose drug A or B. We might represent the cases as follows:

Drugs	Belief	Credence	Facts
A	Cure	0.5	Cure
B	Partial Cure	1	Partial Cure
C	Death	0.5	Death

Dinner	Belief	Credence	Facts
Steak	Morally Permissible	0.9	Morally Permissible
Tofu	Morally Permissible	1	Morally Permissible

Martha's high credence in the permissibility eating steak helps to mitigate the recklessness of her choice, and contributes to the sense that she is not blameworthy. Another feature of her epistemic situation that might be thought to mitigate Martha's blameworthiness is that that the doubts Martha has are comparable to doubts she would have about most philosophical conclusions. We might think it would be unreasonable to expect her to be any more confident; such that when she chooses steak it would not be possible for her to be any more certain. It is very difficult – if not impossible – for a conscientious epistemic agent to achieve complete certainty in most philosophical claims, and a greater degree of certainty might even make us suspicious. If this is right, then Martha's 0.9 credence makes her about as sure as it is reasonable to be about a moral claim, and so plausibly carries much more weight than the same credence for a factual claim, where certainty is often more achievable. Compare Martha to her more reckless counterpart:

Lunch. Louise is deciding whether to have steak or tofu for lunch. She prefers steak, but knows there are ethical questions around meat-eating. She has studied the relevant biological and philosophical literature, and is still unsure. She is inclined to think that meat eating is permissible, but she is not completely sure if this is what the evidence supports, and furthermore she suspects that wishful thinking might be influencing her evaluation of the evidence. So, she has doubts about the correct answer. As a matter of fact, Louise is right that it is permissible for her to order steak, and a fully informed person in her position would know this.

It seems far less plausible that Louise's recklessness is blameless, as it seemed that Martha's could be. While both Louise and Martha make reckless choices, Louise's choice is much more reckless because she is much less sure that her beliefs are correct.

Louise could have easily chosen the low risk option until she had investigated the relevant considerations more carefully, and the fact that she does not is plausibly what contributes to her deserving a more negative evaluation.

A further relevant feature in determining how reckless a choice is, and what kind of evaluation agents deserve is the agent's evidential state. Consider the following example:

Breakfast. Nora is deciding whether to eat steak or tofu for breakfast. She prefers steak, but knows there are ethical questions around meat-eating. She has not bothered to study the relevant biological and philosophical literature, even though as a well-educated philosophy student, it would not be difficult for her to do so. Nevertheless, she is almost (but not completely) certain that it is permissible to eat steak. This is mostly because meat-eating seems natural to her, and most members of her community do not find eating meat morally problematic. As a matter of fact, Nora is right in that a fully informed person in her position would know that meat-eating was permissible.

We could represent Nora's credence in the permissibility of eating meat at 0.9, so just as high as Martha's. However, despite having identical credences, Nora seems much more deserving of negative evaluation than Martha. One explanation of this is that Martha's evidential basis for her high credence is much better than Nora's, and this means that her choice is much less reckless than Nora's. Martha has done everything that we could expect of her in investigating the morally relevant features of the situation. In contrast, Nora could have done much more – as a philosophy student she has the capacity to do much more to inform herself of the morally relevant features and better equip herself to reach a conclusion. Instead she has lazily relied on the

opinions of others. While both are reckless, Nora's choice is much more reckless because it is based on a poor evidential foundation. These various features of Martha's situation all contribute to mitigating the recklessness of her action, and the blame we might think she deserves.

4.2 Inverse Akrasia

If moral recklessness was a vice, we would expect agents who act morally recklessly to be blameworthy. Huck Finn was supposed to be an example of praiseworthy morally reckless action, and so a counterexample to the idea that moral recklessness is always blameworthy. However, further consideration of Huck's motivations allow us to accommodate the intuition that Huck Finn is praiseworthy while maintaining that moral recklessness is a vice.

Actions can be more or less praiseworthy, and different features of those actions can exhibit different degrees of praise and blameworthiness. Huck's action has some praiseworthy aspects – it is part of the description of the case that Huck Finn acts in response to the right kinds of things; the recognition of Jim's humanity, the desire to help a friend, for example. If he had instead freed Jim purely out of a desire to cause mischief, it is not so clear that he would deserve praise. However, in so far as Huck acts recklessly, he is not to be praised. The blameworthiness of his recklessness is compatible with there being other features of his action that are praiseworthy, and it is compatible with the possibility that his action overall has more praiseworthy features than blameworthy features. Huck's action is praiseworthy, but it is not

maximally so. Had Huck freed Jim without acting recklessly, if he had freed Jim in the knowledge that his action was right, his action would have been even more praiseworthy. This would have removed any risk of error, error which might occur due to his false and objectionable belief that property rights can sometimes apply to people¹⁴. The normative externalist position cannot capture this aspect of the Huck Finn case – its explanation for Huck Finn’s praiseworthiness is that moral beliefs are irrelevant to the moral worth of our actions, but this means that it cannot distinguish between cases of inverse akrasia like Huck Finn’s case, and cases in which agents do the right thing because they know it is the right thing. An alternative explanation that can make this distinction is that Huck’s action is praiseworthy because it has features that are sufficiently praiseworthy so as to outweigh any other features of it that would otherwise make it not praiseworthy. This is consistent with moral recklessness being in general blameworthy rather than praiseworthy. To illustrate this, compare an agent who acts just as recklessly as Huck, but without the praiseworthy qualities that redeem his action.

Tuck Finn. Huck’s cousin, Tuck, lives in the same society as Huck and has the same moral opinions. Like Huck, Tuck believes, as everyone in his society “knows,” that helping a slave escape amounts to stealing, and stealing is wrong. Tuck harbours an intense and unreasonable dislike for Miss Watson, and in order to spite her conspires to help one of her slaves escape.

¹⁴ (Sliwa 2016) also makes this point.

It seems fairly clear that Tuck is not to be praised for his action. This is the case even if we think he does the right thing. In so far as Tuck does the right thing, he does so accidentally. He is not to be praised for this accidental right action because the motivations out of which he does it are in no way admirable – he believes what he is doing is morally wrong, and his main motivation is a desire to spite Miss Watson. We might think of this as an example of a case in which praiseworthiness and moral rightness come apart¹⁵. To bolster the point, notice that Huck's akrasia only seems unproblematic when the agents end up doing the right thing. Compare the following case:

Mafia. Dmitri, like most people in his society, believes that killing is wrong. He is a member of the mafia, and as well as believing that killing is wrong values friendship and loyalty to family very highly. An associate has wronged his brother, and it seems to him that the only respectable thing to do would be to kill the associate, even though this would be morally wrong. Although he knows it is wrong, he just cannot bring himself to let the wronging slide and shame his family, so he sets out to murder the associate.

This is an inverted inverse akrasia case. While Huck Finn does the right thing akratically, Dmitri ends up doing the wrong thing akratically. Unlike Huck Finn, he seems blameworthy. This suggests that part of what the intuition that Huck Finn is praiseworthy is tracking is the fact that he does something good rather than bad in spite of false beliefs. There are possible explanations for the sense that this intuition is correct that are neutral on the issue of moral recklessness. For example, it is possible

¹⁵ See (Hawthorne and Srinivasan 2013) for further discussion of this.

that we judge Huck as praiseworthy is because in unknowingly responding to reasons that make freeing Jim morally right, he is taking a first step away from the false views held by his society, and so beginning to equip himself with the tools to do the right thing in the future. What all this suggests is that the Huck Finn case does not on its own demonstrate any asymmetry in the significance of factual and normative uncertainty, and nor it does it show that moral recklessness is not a vice, since there are explanations of the Huck Finn case that are compatible with both factual and moral beliefs being equally significant for appraisal, and with moral recklessness being a vice.

4.3 Moral Heroism and Moral Inertia

The third apparent asymmetry was that while the kind of factual cautiousness employed in *Drugs* is a good general policy for decision making, the same cannot be said for comparable moral cautiousness. The worry was that when employed as a long term policy, moral cautiousness inhibits moral progress, because it makes it unacceptable for the morally conscientious to act on anything less than certainty, and it can seem that considerations of peer disagreement mean that a very high credence in a counter-culture moral view would be unjustified. I want to argue that this worry is unfounded. It is unfounded because it arises only on the assumption that the disagreement of the authorities of one's community is evidence that undermines justification for unorthodox moral beliefs, and this assumption is false.

Even if we agree that disagreement constitutes evidence that should lead us to revise our credences, it is typically only the disagreement of epistemic peers that should push us to revise our credences. One way to resist the assumption that community disagreement constitutes evidence against one's view is to deny that the moral authorities of one's community are epistemic peers. The standard account of epistemic peerhood says that epistemic peers need to have the same evidence and be equally competent in assessing their evidence – they should be equally intelligent, conscientious, and free from bias (see Christensen, 2009; Kelly, 2005). It is not clear that the putative moral authorities of one's community meet these criteria for epistemic peerhood¹⁶. For instance, it is not clear that they have the same evidence. People who disagree about morality often do so while citing different considerations in favour of their views, and may not recognise the same kinds of considerations as evidence. Even if they do recognise the same considerations, they may weight their importance differently¹⁷. It is also far from clear that communities who have false moral views can be said to be free from bias.

Even if it were true that our moral heroes and their opponents were epistemic peers, it need not be true that morally heroic actions are always reckless, since it is very possible that the beliefs of moral heroes have better epistemic standing than those of

¹⁶ Indeed, King (2012) argues that we have very few genuine epistemic peers.

¹⁷ See (Sticker, forthcoming) who argues that ethical disagreements are not the kind of disagreements that require us to lower our credences, and this is because our ethical views are heavily influenced by what the agent considers as relevant evidence bearing on ethical questions.

those who disagree with them. As highlighted in section 4.2, the epistemic standing of one's beliefs matters for whether or not acting on them is reckless. In so far as moral heroes have high credences in beliefs that are based on good evidence, we need not think that they are reckless to act on them. It is worth noting that if they really are moral heroes rather than moral villains, it should not be impossible for epistemically responsible conduct to lead them to true beliefs, and the mere disagreement of others need not undermine well-supported true beliefs¹⁸.

As such, there is no reason why avoiding moral recklessness should lead to moral inertia when employed as a long term policy. However, what does come out of this discussion is a reminder of the importance of seeking out good evidence when forming our moral beliefs, and that mere reliance on the opinions of our communities will not suffice for justification. An example of moral villainy illustrates this:

Forums. Martin spends a lot of time on online forums discussing politics. He shuns the mainstream press and reads only publications recommended by others in the forum. He has come to form some unorthodox moral opinions to the effect that foreigners ought to be sought out and murdered, and that the Aryan race is to be preserved at all costs. Most people in his community beyond the internet forums disagree with him, but he believes he is right, and sees himself as standing up for what is right. In an effort to do the right thing, he tries to encourage his Turkish neighbours to leave by shouting obscenities at them in the street.

The reason why Martin's behaviour is wrong while Rosa Parks' behaviour is right is that Martin's is not based on good evidence at all, while Rosa's (I assume) is. Moral

¹⁸ See also (Kelly 2010).

heroism is compatible with moral conscientiousness, so long as our moral heroes are epistemically responsible. When they are not, they risk being like Martin the moral villain, and they act recklessly. Morally reckless action, just like factually reckless action, is to be avoided.

5. Conclusion

I have argued here that there is no asymmetry between normative and factual uncertainty, and I have shown how it is possible to explain away some apparent asymmetries between normative and factual uncertainty by closer consideration of the cases. Since the asymmetries are merely apparent, there is no reason to think that moral recklessness is not a moral wrong, as factual recklessness is. If this is correct, then normative Jackson cases are no less possible than factual Jackson cases¹⁹.

References:

- Arpaly, Nomy. 2002. *Unprincipled Virtue: An Inquiry Into Moral Agency*. Oxford University Press.
- Bykvist, Krister. 2014. "Evaluative Uncertainty, Environmental Ethics, and Consequentialism." In *Consequentialism and Environmental Ethics*, edited by A. Hiller, I. Ramona, and L. Kahn. Routledge.
- Christensen, David. 2009. "Disagreement as Evidence: The Epistemology of

¹⁹ I am particularly grateful to Jessica Brown and Brian Weatherson for reading and giving invaluable comments on various earlier incarnations of this paper. I am also heavily indebted to an anonymous reviewer for their helpful and charitable comments; and to the organizers and audience of Formal Ethics 2017, University of York. Additionally, I am grateful to friends and colleagues at the University of St Andrews and the University of Michigan for many enjoyable discussion about these topics, in particular: Marvin Backes, Lisa Bastian, Joseph Bowen, Miguel Egler, Rebecca Harrison, Zoë Johnson-King, Ethan Landes, Stefano Lo Re, Tom McKenna, Quentin Pharr, Lewis Ross, Cat Saint-Croix, Justin Snedegar, and Elise Woodard. This research was supported by an AHRC Scottish Graduate School for Arts and Humanities Studentship.

- Controversy." *Philosophy Compass* 4 (5). Wiley-Blackwell.
- Graham, Peter A. 2010. "In Defense of Objectivism about Moral Obligation." *Ethics* 121 (1).
- Guerrero, Alexander A. 2007. "Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution." *Philosophical Studies*.
- Harman, Elizabeth. 2015. "The Irrelevance of Moral Uncertainty." In *Oxford Studies in Metaethics, Volume 10*. Oxford: Oxford University Press.
- — —. 2016. "Morally Permissible Moral Mistakes." *Ethics* 126 (2).
- Hawthorne, John, and Amia Srinivasan. 2013. "Disagreement Without Transparency: Some Bleak Thoughts 1." In *The Epistemology of Disagreement*. Oxford University Press.
- Jackson, Frank. 1991. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101 (3). The University of Chicago Press: 461–82.
- Kelly, Thomas. 2005. "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology, Volume 1*, edited by John Hawthorne and Tamar Gendler. Vol. 1. Oxford University Press.
- — —. 2010. "Peer Disagreement and Higher Order Evidence." In *Social Epistemology: Essential Readings*, edited by Alvin I Goldman and Dennis Whitcomb. Oxford University Press.
- King, Nathan L. 2012. "Disagreement: What's the Problem? Or A Good Peer Is Hard to Find." *Philosophy and Phenomenological Research* 85 (2). Wiley-Blackwell.
- Littlejohn, Clayton. forthcoming "A Plea for Epistemic Excuses." In *The New Evil Demon Problem*, edited by Fabian Dorsch Julien Dutant. Oxford University Press.
- Lockhart, Ted. 2000. *Moral Uncertainty and Its Consequences*. Vol. 111. OUP
- Moller, D. 2011. "Abortion and Moral Risk." *Philosophy* 86 (3): 425–43.
- Sepielli, Andrew. 2009. "What to Do When You Don't Know What to Do." *Oxford Studies in Metaethics* 4.
- Sliwa, Paulina. 2016. "Moral Worth and Moral Knowledge." *Philosophy and Phenomenological Research* 93 (2).
- Sticker, Martin. forthcoming. "Peer-Disagreement about Restaurant Bills and Abortion." *Grazer Philosophische Studien*.
- Sutton, Jonathan. 2007. *Without Justification*. MIT Press.
- Weatherson, Brian. 2014. "Running Risks Morally." *Philosophical Studies* 167 (1).
- Williamson, Timothy. 2017. "Ambiguous Rationality." *Episteme* 14 (3). Cambridge

University Press: 263–74.

Zimmerman, Michael J. 1997. "Moral Responsibility and Ignorance." *Ethics* 107 (3): 410–26.

— — — 2014. *Ignorance and Moral Obligation*. Oxford University Press.

— — —. 2008. *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge University Press.